THE UNIVERSITY OF CHICAGO

ESSAYS IN LABOR ECONOMICS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS

BY
NADAV MORDECHAI KUNIEVSKY

CHICAGO, ILLINOIS
AUGUST 2024

Dedicated to Natalie, Ethan, Sonia, and all my family

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I am deeply thankful to my advisors, Stephane Bonhomme, Ben Brooks, Derek Neal and Jack Mountjoy, for their helpful comments, discussions, and fascinating "behind-the-scenes" conversations about how the sausage is made. I am also grateful for the opportunity to be part of The University of Chicago community and to engage with the faculty of the Economics department and Booth School of Business. Some of these conversations and interactions have left a significant mark on me and were truly inspiring.

I am especially indebted to the friendship and support of the people I met during my time here, particularly Scott Behmer, Hazen Eckert, Francesco Ruggieri, and Shannon Hsuan-Ming Hsu. I want to thank them for the enlightening and interesting conversations about both my work and theirs. These types of discussions are the main reason I recommend individuals undertake this kind of intellectual endeavor, in far-off shores with far-off people. I also want to thank them for their support during the last couple of months of my PhD, which have been marked by particularly painful times in Israel. Their support made this period feel less isolating.

I want to thank my family: Sonia, Ethan, Natalie, my parents, and my sisters for their support. I am indebted to my parents and sisters who helped me get to where I am, providing me with not only information but also the tools to understand, learn, grow, and overcome what I did not think was possible. I am also indebted to Ethan, my son, who, with a few words, keeps reminding me what's important and where everything lies. His laughter, his jokes, his never-ending questions, and our conversations are clearly the reason I get up in the morning.

Most of all, I want to thank my wife, the love of my life, Natalie Goldshtein. Without her, none of this would have been written, thought about, or discussed. She kept me going through easy Sunday mornings and treacherous Mondays, and through all of it, she was beyond supportive. Our conversations, her understanding, and her endless support were nothing short of a miracle, and I can't thank her enough.

Finally, I want to thank Sonia, our beloved dog, who passed away in the last year of this journey. Her love, comforting welcomes, and her unbelievably dignified demeanor are deeply missed.

ABSTRACT

All of our choices and all that sets us apart are governed by what we can do, what we want to do, and what we know. This dissertation aims to quantify two of these channels to better understand why we differ.

The first two chapters focus on what we know and how it shapes societal gaps. The first chapter attacks the question of how much of the gap in choices across social groups is driven by differences in outcomes of choices or by differences in the quality of information these groups have about their respective outcomes. I study this question in the context of the college enrollment gap between white and Hispanic high school students. To assess how much of the gap can be attributed to each channel, I introduce a novel decomposition approach and show how we can use a structural model to operationalize and quantify the role of each component. I find that the main driving force behind the college enrollment gap is differences in potential returns, while differences in information quality across the two groups contribute to narrowing the gap.

The second chapter tackles the question of whether informational asymmetries among firms can account for all observed wage gaps across social groups. I build a common-value auction model in the labor market with unspecified information structures. In this model, firms meet heterogeneous workers with unobserved productivity and extend wage offers based on their information about worker productivity and competing offers. Using the American Community Survey data, I show that wage disparities among Black and white men and women can arise in an economy where different social groups have identical productivity distributions, but firms have different types of information on these different workers, such that the only driving force behind the wage gap is the information.

Finally, the last chapter departs from the notion of *what we know* and turns to discuss what we can do through the lens of intergenerational mobility. It focuses on two common measures of

intergenerational mobility—the Intergenerational Elasticity (IGE) and Rank-Rank coefficients. In it, I employ Yitzhaki's theorem to express these coefficients as weighted averages of the underlying causal mechanisms driving mobility. The chapter highlights the challenges of interpreting cross-country comparisons using IGE or Rank-Rank coefficients due to the regression weighting scheme. It shows that while the Rank-Rank coefficient is more interpretable for positional mobility, it lacks insights into the underlying mechanisms driving mobility across countries.

CHAPTER 1

BRIDGING THE GAP: INFORMATION, RETURNS AND CHOICES

1.1 Introduction

In social systems, where individuals' life trajectories are shaped by choices, understanding the determinants of these choices is crucial, particularly in the pursuit of equality. Standard economic models assume that individuals weigh the costs against the benefits of their decisions. However, it is rarely the case that individuals can perfectly predict the outcomes of their choices. In reality, they operate under significant uncertainty and have limited predictive capabilities about the consequences of their actions. This gap in information and prediction abilities affects the choices different people make, potentially widening or narrowing societal inequities. Therefore, it is essential to assess the extent to which these frictions contribute to differences in decision-making processes and choices.

In this paper, we focus on quantifying how differences across social groups in their ability to predict outcomes contribute to the choice gap across these groups in a binary choice setting. To answer this question, we use a simple choice model framework (Roy (1951)), where individuals facing a binary choice opt in if they perceive the potential returns to be higher than their threshold. We assume that individuals receive informative signals prior to making the decision and use them to form beliefs that guide their choice. We focus on the difference in choice behavior across two groups that stems from members of the two groups having access to different quality signals, which affects the quality of their predictions. Specifically, we measure the *quality of information* each group has by quantifying the share of variance in returns that can be explained by the signals observed by each group. We then say that one group has better information than the other if that group can explain a larger share of variance in their returns. In our analysis, we model the total returns variance as stemming from both the actual uncertainty in returns, which is driven by the underlying data-generating process and the model uncertainty regarding this underlying data-generating process.

In our model, differences in choice are driven by, first, the underlying distribution of returns, and second, the quality of information on these returns. This bifurcation of the choice problem motivates us to adopt a decomposition method akin to that of Kitagawa (1955), Blinder (1973), and Oaxaca (1973) to explore what drives the choice gap. Our approach breaks the choice discrepancy into two channels: the information channel and the returns channel. The information channel quantifies how much of the gap is driven by the fact that the two groups have access to different information sources. It does so by equalizing the information quality across the two groups, holding the returns distribution fixed, and examining how the choice gap changes. The residual difference, as captured by the returns channel, examines what the choice gap would be if we equalized the net returns between the two groups while maintaining their distinct information qualities on those returns.

We apply a decomposition approach to examine the 8% gap in college attendance rates between Hispanic and White students in Texas. To do so, we use administrative data from Texas, which includes information on whether individuals attend a 4-year college and their post-high school earnings. We assume that these high school students observe informative signals on the monetary returns to college, drawn from a Gaussian distribution, which they use to form beliefs on returns. They are then self-select into college based on their posterior beliefs about the monetary returns from college, opting in if their beliefs are higher than their threshold.

Although in our analysis we impose a Gaussian structure, key components of the model are nonparametrically identified. In our model, beliefs dictate choice patterns, this allows us to use choice data to nonparametrically identify the distribution of beliefs and earnings for each group. Specifically, building on the marginal treatment effect literature (Heckman and Vytlacil (2005)) we show how in our model the beliefs distribution is identified. We assume that we have a continuous instrument that shifts the cost of attendance. In our empirical exercise, this instrument is the distance to a 4-year college. We assume that, conditional on a set of controls, distance to college is independent of both information and earnings and affects only the cutoff value (Card (1995), Carneiro et al. (2011), Nybom (2017), Kapor (2020), Walters (2018), Mountjoy (2022)). We then trace how small changes in the instrument change the conditional expectation of earnings. A small increase in the cost of attendance pushes out those individuals whose new cost is higher than their beliefs. Using the assumption of rational expectations, tracking these changes in the expected earnings tells us about the beliefs of these marginal individuals who are responding to the small cost change. Similarly, tracking how these changes affect the propensity of attending college reveals the share of individuals with these beliefs.

Using our decomposition approach as outlined above, and estimated parameters of the Gaussian model, we find that differences in the information quality across the group contributed to shrinking the choice gap. Specifically, the information channel shows that equating the information quality across groups would increase the choice gap by approximately 7.7% (around 97% of the original choice gap). The decomposition exercise also shows that most of the current gap in choice is driven by differences in the returns distribution faced by Hispanics and Whites. Specifically, we find that the potential returns for college for Whites are much larger than those of Hispanics, and that these differences drive most of the choice gap.

We focus in our analysis on the differences in the choice gap that are driven by differences in the quality of information. This is not the only approach to measuring the effect of information differences. In the Appendix, we introduce an additional decomposition approach, where instead of equating the information quality across groups, we equate the information structure. This decomposition approach, which builds on tools from the robust mechanism (Bergemann et al. (2022), Bergemann and Morris (2016), Bergemann and Morris (2013)) literature, allows us to bound the full set of counterfactuals nonparametrically and enables us to depart from the Gaussian distribution assumption.

In the second part of the paper we turn to ask how parity in choice can be achieved by considering a policymaker that wants to close the gap by providing Hispanics with additional new information. We postulate that this policymaker, acting as a statistician with access to information on earnings for individuals who attend college and those who do not, could provide an informative signal to each high school student about their potential income. We then ask how accurate must this additional information be? Our findings suggest that to effectively close the gap, this new information must be able to explain either 24% of the variance in college earnings or 49% of the variance in non-college earnings. We explore the feasibility of achieving this level of information accuracy using data available to schools. Our administrative data is utilized to predict earnings 12 to 14 years after high school graduation for both college attendees and non-attendees. We show that, at most, we can explain 10% of the quarterly earnings variance. This suggests that closing the attendance gap through the provision of information necessitates the development of more accurate sources for earnings prediction.

Related Literature. This paper contributes to an extensive body of literature on human capital investment decisions, anchored by the foundational work of Ben-Porath (1967). Our study intersects with research focused on the impact of monetary returns on such choices, as explored in studies by Willis and Rosen (1979), Cunha and Heckman (2007b), Walters (2018), Abdulkadiroğlu et al. (2020), and Freeman (1971). These papers typically make assumptions about what information is used by individuals to beliefs about returns—often measured based on observable factors—and analyze how these beliefs factor into decision-making processes. Our approach differs from them by examining how differences in the quality of information to individuals influence their choices and drive the gap across groups. Focusing on the the quality of information and not the specific beliefs, or variables used in the inference process.

Another significant aspect of our research aligns with studies that investigate the nature of individuals' beliefs, such as those by[1] Manski (2004), Wiswall and Zafar (2015), Zafar (2011), Wiswall and Zafar (2021), and Diaz-Serrano and Nilsson (2022). These works delve into systemic differences and biases in beliefs among groups defined by socio-economic status. Our paper extends this inquiry, utilizing these findings to illuminate not just the distribution of beliefs but also the quality and extent of information available to these groups.

As discussed above, methodologically, our study builds upon the Marginal Treatment Literature, particularly the work of Heckman and Vytlacil (2005). This approach has previously been employed to examine the marginal treatment effects on returns to schooling, as demonstrated by Carneiro et al. (2011), Carneiro and Lee (2009) and Mountjoy (2022). Similar to some of these studies, we link the marginal treatment effect to beliefs. Eisenhauer et al. (2015) employed this structure to conduct a cost-benefit analysis of programs, focusing on agents' ex-post and ex-ante costs—closely paralleling our usage. Canay et al. (2020) and d'Haultfoeuille and Maurel (2013), in the context of college decisions and discrimination, demonstrate how the Roy model can identify ex-ante beliefs and preferences, aligning with our methodological approach.

Our work related to recent research by Bohren et al. (2022) on systemic discrimination. Their study, akin to ours, identifies two main sources of systemic differences between social groups. The first, termed 'technological systemic discrimination', aligns with our focus on differences in return distribu-

1. See overview of the literature on beliefs elicitation in Giustinelli (2022)

tions and captures disparities across groups in certain outcome variables. The second, 'informational discrimination', pertains to disparities arising from varied information available to decision-makers across groups. Our research differs in its concentration not on discrimination towards individuals but on the decisions individuals make about themselves and how these systemic forces shape it, with a specific focus on the quality of information rather than its structure. We further explore a distinct measure related to this in our Appendix.

While our primary focus is on educational decisions, our decomposition approach has broader applications. It can illuminate how information asymmetries contribute to decision-making disparities across various contexts. Recent studies, including those by Arnold et al. (2018), Arnold et al. (2022), and Canay et al. (2020), have explored the influence of judicial preferences and biases in decision-making. There is a growing interest in understanding how decision-making signals contribute to these disparities. Our decomposition methodology seeks to address these nuanced aspects of decision-making processes.

The remainder of the paper proceeds as follows. Section 1.2 describe our framework and decomposition approach. Section 1.3 describe the data and some descriptive statistics. Section 1.4 describes some empirical patterns on earnings and information. Section 1.5 discuss the estimation results. Section 1.6 discuss counterfactual effects of providing additional information and section 1.7 concludes.

1.2 Framework

We consider a population of high school graduates, indexed by i. At the end of high school, each graduate must decide whether or not to attend college. The objective of individual i is to maximize earnings. Denote by Y_1^i earning for an individual i who attends college and by Y_0^i their earnings if they do not attend. We assume that earnings are generated according to

$$Y_1^i = \alpha_1^i + u_1^i,$$
$$Y_0^i = \alpha_0^i + u_0^i,$$

where α_d^i, $d \in \{0, 1\}$, is the structural component of earnings and u_d^i is an unpredictable component of earnings, satisfying $E[u_d^i|\alpha_1^i, \alpha_0^i] = 0$. Before deciding whether to attend college, each student i observes an informative signals on the their individual structural component of earnings. Specifically, we denote by $\boldsymbol{S}_i \in \mathcal{S}$ the vector of realized signals that individual i observes and assume that $\boldsymbol{S} \perp u_d^i|\alpha_1^i, \alpha_0^i$. Our model separates earnings into two components. The first is a structural component, α_1 and α_0, which agents can know and form beliefs about. The second component is u_d, which is unknowable at the time of the decision. These components of earnings include idiosyncratic shocks that can only be known ex-post. Henceforth, α_1^i and α_0^i will be treated as earnings, and the index i will be omitted for clarity when its presence is self-evident.

In our model, signals link outcomes to beliefs; thus, we need to determine how individuals use signals to form beliefs. We adopt the standard approach in economics and model individuals as Bayesian agents. Being Bayesian implies that individuals observe signals, know the correct likelihood function, and update their prior beliefs to form new posterior beliefs over the outcomes. Let $\pi(\alpha_1, \alpha_0, \boldsymbol{S})$ be the joint distribution of outcomes and signals. We assume that this distribution is parameterized by a set of parameters, some of which are perfectly known to the individuals, denoted by θ_k, while on others the individuals may have priors over, denoted by θ_u. We denote by $H(\theta_u)$ and $h(\theta_u)$ the prior cumulative distribution function (CDF) and the corresponding probability density function (PDF) the set of unknown parameters.

After observing the signal realization, $\boldsymbol{S}$, individuals form beliefs on their returns, based on their posterior beliefs $H(\theta_u|s)$ and information in the models. Specifically, denote by $\mathcal{R} = \alpha_1 - \alpha_0$ the structural part of the returns, by $\mathcal{E}[.|s]$ the individuals posterior beliefs, which incorporates the model uncertainty, and $\mathrm{E}[.|s]$ the standard expectation operator then we have that the individuals posterior beliefs over their returns is given by

$$\begin{aligned}\mathcal{E}(Y_1 - Y_0|\boldsymbol{S}) &= \int \mathrm{E}[Y_1 - Y_0|\boldsymbol{S}; \theta_u, \theta_k] dH(\theta_u|\boldsymbol{S}) \\ &= \int E[\alpha_1 - \alpha_0|\boldsymbol{S}; \theta_u, \theta_k] dH(\theta_u|\boldsymbol{S}),\end{aligned}$$

where we used the fact that the unpredictable errors are independent from the signal and have mean zero.

Finally, We assume that individuals incur some cost when attending college, that is a function of observables. We denote by X observed variables, by $c(x)$ the cost of attendance and by then individual i's decision rule is given by

$$D = \mathbb{1}\left[\mathcal{E}[Y_1 - Y_0|\boldsymbol{S}] \geq c(x)\right] = \mathbb{1}\left[\mathcal{E}[\mathcal{R}|\boldsymbol{S}] \geq c(x)\right].$$

Our decision rule suggests that individuals derive risk-neutral utility from earnings but allows high school graduates to possess any utility function that strictly increases with expected returns (Vytlacil (2006)). Modeling utility as an increasing function of returns includes also the standard linear indirect utility function, that has been used in models of school and education choices (Willis and Rosen (1979), Walters (2018), Abdulkadiroğlu et al. (2020)) In our framework, we standardize this utility function to be the identity function. Therefore, $c(x)$ serves as a composite of individual preferences, known monetary and non-monetary costs, and other barriers to college attendance, such as credit constraints, social norms, and additional limitations.

It's also important to keep in mind that although we model here the decision process as a result of one individual's choice, it is likely that the decision to go to college or not is made in conjunction with other parties, such as parents, guardians, or advisors. In this case, the observed signals are all the signals observed by all parties, and the cost is an agglomeration of all members who participate in the decision.

1.2.1 What information may high school students have on their returns

It is worth considering what signals high school students observe prior to deciding about college. Some signals are pieces of information that students receive throughout their lives. For example, students may hear from various media sources about the potential returns to college for different types of students, learn about their ability from test scores, or consult with their parents or high school counselors about their prospects. The quality of this information depends on how well it correlates with the individual's unique characteristics and the future labor market. For instance, educated parents may be better informed about potential jobs for college graduates and their children's unique skill sets. Consequently, they might provide accurate information on post college earnings for college graduates based on their familiarity with different career paths and their children's suitability for these paths.

The above examples are instances of information that decision-makers may be exposed to prior to making their choice. Another source of information may come from the way the labor market itself operates. In section 1.12.1, we discuss in detail two simple examples, demonstrative the richness of information used in the decision process. The first example considers the case where individuals have perfect knowledge of a structural component to wages. For instance, individuals are likely to know their ability. Ability plays a role in determining the potential earnings, but in our model, knowing your own ability is also a piece of information in and of itself. Accordingly, some components in the economy hold two simultaneous roles. First, they affect the actual returns of the individual; separately, they are also used as a signal about earnings.

The second example considers the case where two individuals may observe the same signal realization, but how informative these signals are may differ due to the labor market structure. In this example, we consider two groups that can go to college and become either lawyers or accountants. Each group member observes a signal on whether they will be a high-earning lawyer or a low-earning lawyer. The two groups differ in the share of people who end up becoming lawyers after college. This difference in the labor market for the two groups implies that the same signal may carry different information on returns, which is the thing the decision makers care about.

These examples illustrate the complex and multifaceted nature of information available to high school students as they contemplate their educational futures. The signals they receive come from a variety of sources, each with its own level of reliability and relevance. This also motivates our approach to decomposition, which focuses not on specific pieces of information, but rather on the overall information quality, which is a crucial element in the choice gap.

1.2.2 Gaussian Model

In this section, we restrict the general model we presented above and impose a Gaussian structure. This has two main advantages. First, imposing a Gaussian structure on the signals implies that we can rank information based on the Blackwell Ordering. In general, the Blackwell order (Blackwell (1953)) is a partial order on information structures, where one information structure is said to be

Blackwell more informative than another if it leads to better decision-making outcomes for every decision problem. In the Gaussian case, there is a total order over information. Therefore, imposing the Gaussian structure implies that we can meaningfully say that one group has better information than the other (Chan et al. (2022)). The second advantage is that the Gaussian distribution is fully determined by the first and second moments. This implies that we can fully characterize the beliefs distribution and counterfactual distribution we need for analysis ahead by the mean and our measure of information quality, which we discuss below.

We assume that the signals, $\boldsymbol{S}$, and the structural components of earnings, α_1 and α_0, are jointly distributed as Gaussian. We denote by μ_1, μ_0, σ_1, σ_0, and ρ the means, standard deviations, and the correlation between α_1 and α_0, respectively. We further assume that individuals have "partial rational expectations,", where they know the parameters that govern the marginal distribution of α_1 and α_0, and the correlation between signals and potential earnings. On the other hand we take the more realistic approach that high school students do not know the correlation parameter, ρ, between α_1 and α_0. We denote the individuals prior over ρ, by $H(\rho)$, with expected value given by μ_ρ.

This modeling captures the notion that agents, similar to econometricians, may know how signals are linked to the marginal distribution of earnings, but cannot learn from data and observation on the correlation between potential outcomes. Specifically, in our model the signal realization $\boldsymbol{S}$, does not provide any information on the correlation structure.[2]

The fact that individuals can not learn about correlation from observation also implies also that

2. To see this, notice that

$$\pi(\rho|\boldsymbol{S}) = \frac{\pi(\boldsymbol{S}|\rho)h(\rho)}{\pi(\boldsymbol{S})} = \frac{\int \pi(\boldsymbol{S}, \alpha_1, \alpha_0|\rho) d\alpha_1 \alpha_0 h(\rho)}{\pi(\boldsymbol{S})} = \frac{\pi(\boldsymbol{S})h(\rho)}{\pi(\boldsymbol{S})} = h(\rho)$$

where the second equality is simply the law of total probability and the third equality stems from the fact that for each ρ, the marginal of π with respect to s is always the same individuals beliefs on returns are aligned with the true conditional expectations on returns, in the sense that[3].

$$\mathcal{E}\left[\alpha_1 - \alpha_0 | \boldsymbol{S}\right] = \mathrm{E}\left[\alpha_1 - \alpha_0 | \boldsymbol{S}\right]$$

This implies that the beliefs on the the expected returns are correct, notice that the beliefs on higher moments of the conditional returns distribution may differ from the true underlying distribution of returns.

Next we derive the distribution of beliefs and the share of high school students who opt into college. Given that potential earnings and signals are jointly Gaussian, it follows that returns and signals are also jointly Gaussian distributed:

$$\begin{bmatrix} \boldsymbol{S} \\ \mathcal{R} \end{bmatrix} | x \sim N\left(\begin{bmatrix} \mu_{\boldsymbol{S},x} \\ \mu_{\mathcal{R},x} \end{bmatrix}, \begin{pmatrix} \Sigma_{\boldsymbol{S},x} & \Sigma_{\boldsymbol{S},\mathcal{R},x} \\ \Sigma_{\boldsymbol{S},\mathcal{R},x} & \sigma^2_{\mathcal{R},x} \end{pmatrix}\right).$$

Where $\sigma^2_{R,x} = \sigma_1^2 + \sigma_0^2 - 2\rho\sigma_1\sigma_0$ denote the variance of returns, $\Sigma_{\boldsymbol{S},x}$ denotes the covariance matrix of the signals, $\Sigma_{\boldsymbol{S},\mathcal{R},x}$ is the covariance between signals and returns, and $\mu_{\boldsymbol{S,x}}$ and $\mu_{\mathcal{R},x}$ are the mean values of signals and returns, respectively. We let all variables to be conditional on x. As signals and returns are jointly Gaussian, individuals who observe the signals realization $\boldsymbol{S}$ form the following

3. To see this, notice that

$$\begin{aligned}
\mathcal{E}[\alpha_1|\boldsymbol{S}] &= \int_\rho \int_{\alpha_1} \alpha_1 \pi(\alpha_1|\rho, \boldsymbol{S}) d\alpha_1 dH(\rho|\boldsymbol{S}) \\
&= \int_\rho \int_{\alpha_1} \alpha_1 \frac{\pi(\alpha_1, \boldsymbol{S}|\rho)}{p(\boldsymbol{S}|\rho)} d\alpha_1 dH(\rho|\boldsymbol{S}) \\
&= \int_\rho \int_{\alpha_1} \alpha_1 \frac{\pi(\alpha_1, \boldsymbol{S})}{p(\boldsymbol{S})} d\alpha_1 dH(\rho|\boldsymbol{S}) \\
&= \mathrm{E}[\alpha_1|\boldsymbol{S}]
\end{aligned}$$

where we used the fact the marginal of π for any value of ρ is the same and that $\boldsymbol{S} \perp\!\!\!\perp \rho$. We can similarly show that $\mathcal{E}[\alpha_0|\boldsymbol{S}] = \mathrm{E}[\alpha_0|\boldsymbol{S}]$, which implies the desired result posterior beliefs on their returns:

$$E[\mathcal{R}|\boldsymbol{S},x] = \mu_{\mathcal{R},x} + \Sigma^T_{\boldsymbol{S},\mathcal{R},x}\Sigma^{-1}_{\boldsymbol{S,x}}(\boldsymbol{S} - \mu_{\boldsymbol{S},x}). \tag{1.1}$$

We now can write explicitly the decision rule for an individual with cost $c(x)$ and signal realization $\boldsymbol{S}$ in our Gaussian model as:

$$D = \mathbb{1}\left[E[\mathcal{R}|\boldsymbol{S},x] \geq c(x)\right] = \mathbb{1}\left[\mu_{\mathcal{R},x} + \Sigma^T_{\boldsymbol{S},\mathcal{R},x}\Sigma^{-1}_{\boldsymbol{S,x}}(\boldsymbol{S} - \mu_{\boldsymbol{S},x}) \geq c(x)\right].$$

Owing to the linearity of the joint Gaussian distribution, we can derive the proportion of students who opt to attend college. To do so, we first notice that the share of individuals who would opt to go to college depends on the share of individuals whose prediction of the returns is higher than the cost. We therefore first derive the distribution of beliefs in the population from the posterior beliefs in equation 1.1.

$$E[\mathcal{R}|\boldsymbol{S},x] \sim \mathcal{N}\left(\mu_{\mathcal{R},x}, \Sigma^T_{\boldsymbol{S},\mathcal{R},x}\Sigma^{-1}_{\boldsymbol{S},x}\Sigma_{\boldsymbol{S},\mathcal{R},x}\right),$$

where we used the linearity of the posterior beliefs. By knowing how beliefs are distributed in the population, we easily derive the share of individuals with cost $c(x)$, who would go to college:

$$P(D=1|c(x)) = \Phi\left(\frac{\mu_{\mathcal{R},x} - c(x)}{\sqrt{\Sigma^T_{\boldsymbol{S},\mathcal{R},x}\Sigma^{-1}_{\boldsymbol{S},x}\Sigma_{\boldsymbol{S},\mathcal{R},x}}}\right),$$

where Φ denotes the standard normal CDF. Henceforth, we will omit x in our discussion, except in cases where it contributes significantly to the analysis.

1.2.3 Information Quality

In our framework, individual choice is influenced by two factors: the net returns $\mathcal{R} - c(x)$ and the individuals' ability to predict these returns. Our analysis seeks to understand how these elements impact decision-making across different social groups. To do so, we need to define formally how we can measure the individuals' prediction ability. In this section, we define a measure of *information quality* that captures the prediction quality of the individuals. Specifically, we quantify the quality of information individuals have by the coefficient of determination, often denoted by R^2 (R-Squared). This metric measures the proportion of the variance in returns that can be explained by their signals, relative to the total variance in returns, from the high school perspective.

$$R^2 = \frac{\mathrm{Var}(E[\mathcal{R}|\boldsymbol{S}])}{\mathrm{Var}_{total}(\mathcal{R})}.$$

The total variance of returns, $\mathrm{Var}_{\mathrm{total}}(\mathcal{R})$, is influenced by two sources of uncertainty. The first is individual uncertainty regarding their specific returns, as discussed above. The second source of uncertainty is *model uncertainty*, stemming from the fact that individuals can not know ρ.

We can also derive an explicit expression for the total variance of returns as follows:

$$\begin{aligned}\mathrm{Var}_{\mathrm{total}}(\mathcal{R}) = \mathrm{Var}_{\mathrm{total}}(\alpha_1 - \alpha_0) &= E[\mathrm{Var}(\alpha_1 - \alpha_0|\rho)] + \mathrm{Var}(E[\alpha_1 - \alpha_0|\rho]) \\ &= \sigma_1^2 + \sigma_0^2 - 2\sigma_1\sigma_0\mu_\rho,\end{aligned}$$

where the first equality follows from the law of total variance. Using this expression we have that information quality is given by

$$R^2 = \frac{\mathrm{Var}(E[\mathcal{R}|s])}{\sigma_1^2 + \sigma_0^2 - 2\sigma_1\sigma_0\mu_\rho} \tag{1.2}$$

This R^2 differs from the standard coefficient of determination, as it accounts for both fundamental uncertainty and subjective uncertainty over the underlying data generating process. Similar to the standard R^2, this measure ranges from 0, implying that the information available does not reduce any uncertainty, to 1, implying that the information and the way it's correlated with the potential earnings resolves all objective and subjective uncertainties.[4]

As we discuss further in Section 1.2.5, the prior distribution $H(\rho)$ is not directly identifiable from the data. This means that we cannot determine the exact shape or parameters of the distribution solely by observing the data. However, we can partially identify the support of $H(\rho)$, taking into account the restrictions implied by the covariance matrix between the signals and the potential earnings. This partial identification allows us to make more informed guesses about the possible distribution H, even if we do not know the precise distribution.

In our main analysis, we impose the assumption that $\mu_\rho = 0$, implying that individuals hold a symmetric prior over the full set of ρs, $[-1, 1]$. Under this assumption, the ex-ante uncertainty in returns is determined by the variance in the marginal distribution, as captured by σ_1^2 and σ_0^2.

One motivation for this assumption is the concept of equi-ignorance. If the individuals lack any information about the correlation of potential outcomes, similar to econometricians, they would assign equal weight to each correlation value, implying a uniform prior over the possible data-generating processes, this would result in a uniform distribution, centered around 0. Another motivation comes from ex-post inference. If individuals have a degenerate prior distribution with $\rho = 0$, observing one realization of the potential outcomes would not enable them to infer the other. This is akin

4. To see that the R^2 in equation 1.2 is always less or equal than 1, note that the set of feasible ρs must satisfy at least $\text{Var}[\alpha_1 - \alpha_0|s] \leq \sigma_1^2 + \sigma_0^2 - 2\sigma_1\sigma_0\rho$, as explained in more detail in 1.2.6. Therefore, as the denominator integrates over the set of ρs such that the variance of returns is higher than the variance of beliefs, the value of the R^2 is always less than 1. Further, the value 1 is achieved in the case where $\text{Var}(\mathbf{R}|s) = \sigma_1^2 + \sigma_0^2 + 2\sigma_1\sigma_0$. This occurs when there is only one possible ρ feasible, $\rho = -1$, and agents have full information.

to econometricians who cannot use observed realized outcomes to infer the value of the unobserved ones.

How does the quality of information and returns affect the decision on going to college? Our measure for information quality implies that higher quality of information implies higher dispersion of beliefs among individuals. Intuitively, if individuals have access to better quality, more accurate, information, then they would respond to it more, and rely on it more when updating their beliefs, instead of relying on the mean returns. Therefore, better information would result in an increasing belief dispersion. Whether higher beliefs dispersion implies that more individuals would attend college is contingent upon the relationship between the cost of attendance and the mean returns in the population, $\mu_{\mathcal{R}}$. Figure 1.1 illustrates the interaction between the mean returns, $\mu_{\mathcal{R}}$, information quality, and cost and how they affect choices. The black line represents the cost. The two red lines represent the survival functions for priors lower than the cost, while the two blue lines represent the survival functions for priors higher than the cost. Dashed lines indicate the posterior beliefs distribution for an agent with high-quality information, and solid lines represent the posterior mean distribution with low-quality information. The figure shows that if the cost is lower than $\mu_{\mathcal{R}}$, increasing the precision of the signal—or enhancing information quality—would reduce college attendance. Conversely, if the mean returns exceed the cost, a reduction in information quality could prove actually increase the share of individuals who opt in to college.

Furthermore, examining the cross-derivative of the share of people who go to college with respect to the difference between the prior mean and the cost shows that when $\mu_{\mathcal{R}} - c \leq 0$, widening the negative difference increases the effect of better information.

$$\frac{\partial^2 \Phi\left(\frac{\mu_{\mathcal{R}}-c}{\sqrt{\mathrm{Var}(\mathrm{E}[\mathcal{R}|\boldsymbol{S}])}}\right)}{\partial\sqrt{\mathrm{Var}(\mathrm{E}[\mathcal{R}|\boldsymbol{S}])}\partial(\mu_{\mathcal{R}}-c)} = -\frac{1}{\mathrm{Var}(\mathrm{E}[\mathcal{R}|\boldsymbol{S}])}\left[-\phi\left(\frac{\mu_{\mathcal{R}}-c}{\sqrt{\mathrm{Var}(\mathrm{E}[\mathcal{R}|\boldsymbol{S}])}}\right)\frac{\mu_{\mathcal{R}}-c}{\sqrt{\mathrm{Var}(\mathrm{E}[\mathcal{R}|\boldsymbol{S}])}} + \phi\left(\frac{\mu_{\mathcal{R}}-c}{\sqrt{\mathrm{Var}(\mathrm{E}[\mathcal{R}|\boldsymbol{S}])}}\right)\right],$$

where we can see that the terms inside the bracket are positive in the case where $\mu_{\mathcal{R}} - c \leq 0$. In the case where $\mu_{\mathcal{R}} - c > 0$, the term in the brackets can be either negative or positive, but tends to be negative as $\mu_{\mathcal{R}} - c \to \infty$. Therefore, we can see that providing better information in the case where $\mu_{\mathcal{R}} - c < 0$, would have a stronger effect if the difference between the prior mean and the cost is larger. It is worth noting also that we can flip the perspective and notice that the effect of increasing the difference between $\mu_{\mathcal{R}}$ and c —by raising the cost or altering the average returns—would be more significant if beliefs are more dispersed.

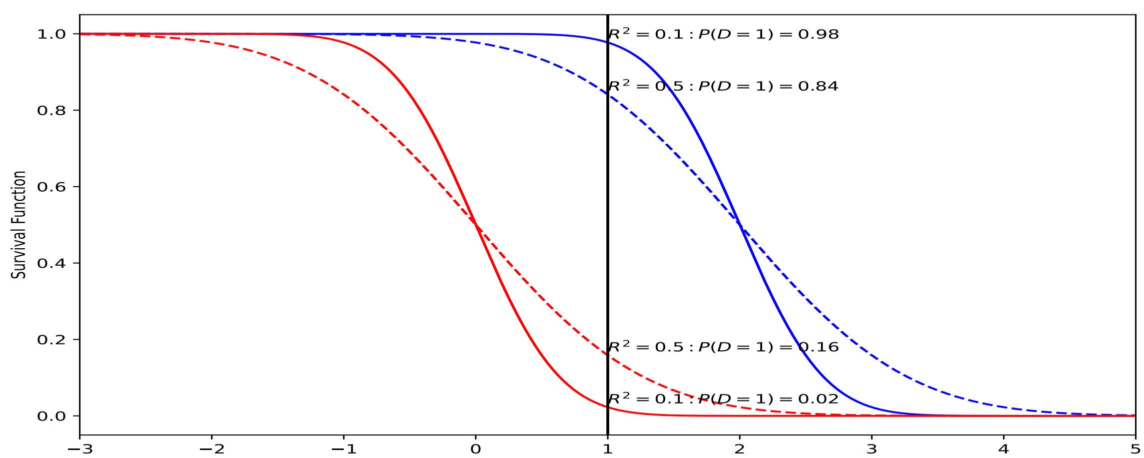

Figure 1.1: Cost, information and Beliefs interaction

Note: This figure illustrates how the interaction between the prior, information quality, cost affect choices. The black line represents the cost. The two red lines represent the survival functions for priors lower than the cost, while the two blue lines represent the survival functions for priors higher than the cost. Dashed lines indicate the posterior beliefs distribution for an agent with high-quality information, and solid lines represent the posterior mean distribution with low-quality information. The figure demonstrates that if the priors are higher than the cost, providing additional information reduces the share of participants from 98% to 0.84%. Conversely, if the prior is lower than the cost, improving the quality of information increases the share of individuals who opt in.

1.2.4 Decomposing the Choice Gap

To quantify the role of information in exploring the gap, we suggest using a decomposition method à la Kitagawa (1955), Blinder (1973), and Oaxaca (1973). In it, we decompose the differences in choices into two components, stemming from the differences in information quality and differences and differences in the returns distribution themselves between two groups. Specifically, we investigate what proportion of individuals would choose to attend college if individuals from different groups, with the same observables, had access to the same quality of information.

Before introducing our decomposition, we define some notation. Let R_g^2 be the quality of information of group g. Let μ_ρ^g denote the mean beliefs on ρ for group g, and similarly, we denote the group components of the returns distribution of group g as μ_g, $\sigma_{1,g}$, and $\sigma_{0,g}$. Finally, let $\text{Var}_{g,g'}(E[\mathcal{R}|\boldsymbol{S}])$ be the counterfactual variance of beliefs for group members g, with the information quality of group g' and the earning distribution of group g.

$$\text{Var}_{g,g'}(E[\mathcal{R}|\boldsymbol{S}]) = R_{g'}^2 \times (\sigma_{1,g}^2 + \sigma_{0,g}^2 - 2\sigma_{1,g}\sigma_{0,g}\mu_\rho^g).$$

This expression captures the beliefs variance if group g had the same quality of information as group g', but faced the an unchanged returns distribution. We then suggest to decompose the choice gap between group a and b as follows:

$$
\begin{aligned}
&P(D=1|\text{Group b}) - P(D=1|\text{Group a}) = \\
&\underbrace{\int_X \Phi\left(\frac{\mu_{\mathcal{R},b,x} - c_b(x)}{\sqrt{\text{Var}_{b,b}(E[\mathcal{R}|\boldsymbol{S},x])}}\right) dF_b(x) - \int_X \Phi\left(\frac{\mu_{\mathcal{R},b,x} - c_b(x)}{\sqrt{\text{Var}_{b,a}(E[\mathcal{R}|\boldsymbol{S},x])}}\right) dF_b(x)}_{\text{Information Channel}} \\
&+ \underbrace{\int_X \Phi\left(\frac{\mu_{\mathcal{R},b,x} - c_b(x)}{\sqrt{\text{Var}_{b,a}(E[\mathcal{R}|\boldsymbol{S},x])}}\right) dF_b(x) - \int_X \Phi\left(\frac{\mu_{\mathcal{R},a,x} - c_a(x)}{\sqrt{\text{Var}_{a,a}(E[\mathcal{R}|\boldsymbol{S},x])}}\right) dF_a(x)}_{\text{Returns Channel}}
\end{aligned}
\tag{1.3}
$$

Where we denote the CDF of X for group g by $F_g(x)$. In our decomposition, the information channel quantifies the extent to which the gap in choices arises from individuals of different groups having access to different qualities of information. These differences may increase or decrease the gap, depending on the relation between the mean returns and costs and the quality of information each group has, as discussed above.

How can we equate predictive ability across two groups? It's instructive to consider cases where members of groups a and b use different models to predict the outcomes of their choices. These models may differ in their set of explanatory variables; one group might use a larger set of variables to explain outcomes, while the other may have a smaller set, potentially resulting in poorer model performance. In our setting, the exact set of variables used is less crucial, as these variables only affect choices through the information they contain about outcomes.

In our decomposition exercise, we focus on equating the quality of prediction across groups, which doesn't necessarily mean equalizing the signals they observe. As discussed in section 1.2.1, such an approach may be impractical. Instead, we can conceptualize this as a counterfactual world where we either provide more variables to the group with the smaller set or remove some explanatory variables from the group with better information. Since the exact nature of these variables is unimportant for the decision rule, we disregard them and focus solely on the quality of information.

It is important to recognize that our analysis is a partial equilibrium exercise, where we use comparative statics to equalize the information quality between the two groups. Typically, information quality is determined endogenously within an equilibrium framework (Coate and Loury (1993), Lundberg and Startz (1983)), and is driven by choices individuals make that form the information environment and what agents can know. Furthermore, the information quality that individuals possess could be influenced by the effort they invest in acquiring it, a concept central to the standard rational inattention model (Caplin et al. (2022); Maćkowiak et al. (2023)). In this decomposition exercise, we do not explore the underlying factors that drive these information discrepancies; rather, we take them as given and investigate the extent of their contribution to the observed disparity.

Finally, we do recognize that our counterfactual choices shares relies on the the second moment of both returns and beliefs. In more general settings, with unrestricted data-generating processes, with more nuisance information structure, equalizing R^2 does not yield a unique counterfactual. In many cases, different joint distributions of signals and outcomes may produce the same R^2 but induce complex choice patterns that contribute to gaps in choices influenced by information. In section 1.12 in the Appendix, we discuss another decomposition approach that equalizes the information structure across groups. This approach does not equalize the ability to predict across groups, but rather equalizes the signals that individuals with similar outcomes receive.

We now examine the second channel, which we call the *returns channel.* This residual component addresses the inverse question: By how much would the share of high school graduates from group a change if we maintained their information quality at R_a^2, but adjusted their returns and costs to match those of group b? This component reveals the extent to which the gap is driven by differences in the outcome distribution itself, rather than information quality. Consequently, we interpret this component as quantifying the portion of the gap attributable to the underlying labor market fundamentals that drive differences in choice.

The two components of the distribution carry distinct policy implications. If the majority of the gap is driven by differences in predictive ability and information, policymakers aiming to close the choice gap should consider providing additional information to group b. This can be done by either "transferring" the superior information from group b to group a or providing additional new information to group a members. This could involve educational interventions, information dissemination, or providing improved prediction tools for group a. Conversely, if the gap primarily stems from variations in the outcome distribution, policymakers concerned with narrowing the disparity in choice should focus on policies that directly influence this distribution. This could include measures such as altering tax structures, providing targeted subsidies, or implementing regulatory changes that affect the underlying returns and costs for both groups. Identifying the primary driver of the gap not only enhances our understanding of its structural roots but also provides actionable insights for policymakers committed to fostering equal opportunities across different groups.

1.2.5 Model Identification and Empirical Specification

This section outlines the key components required for our decomposition exercise. While Appendix 1.10.2 provides a comprehensive nonparametric identification argument based on Marginal Treatment Effect (Heckman and Vytlacil (2005)) and the discrete choice model identification (Matzkin (1992),Matzkin (1993)) literature, here we focus on the essential assumptions we need and demonstrate the identification argument using our simplified Gaussian model.

Our identification hinges on three sets of assumptions. The first set of assumptions pertains to the validity of an instrumental variable. The second set of assumptions regards the heterogeneity of the individuals, and the last set of assumptions focuses on the utility function and the usage of partial rational expectations.

The assumptions needed for our instrumental variable are the standard instrumental variable assumptions (Angrist and Imbens (1995), Vytlacil (2002), Heckman and Vytlacil (2005)). As we detailed in section 1.10.2, we require that the instrument satisfy the exclusion and relevance assumptions. In terms of our models, these assumptions imply that our instrument Z is used as a cost shifter that shifts $c(Z, X)$ (relevance) and satisfies the exogeneity requirement, $\alpha_1, \alpha_0 \perp\!\!\!\perp Z|X$ and $S \perp\!\!\!\perp Z|\alpha_1, \alpha_0, X$.[5] The first independence condition requires that a shift in the instrument, conditional on observables, does not change the distribution of potential outcomes. The second condition implies that, conditional on the set of covariates and the individual potential earnings, the instrument is independent of the signals individuals observe. In other words, we require that two individuals with the same observables, X, and the same potential earnings should draw signals from the same distribution.

As we detail in the next section, in our empirical analysis we use the distance from the student high school to the nearest 4 years college as our instrument. Distance to college has been first used by Card (1995) to estimate college returns and was then used extensively as in instrument to study various outcomes of schooling (e.g Carneiro et al. (2011), Nybom (2017), Kapor (2020), Walters (2018), Mountjoy (2022), Kling (2001),Kane and Rouse (1995),Cameron and Taber (2004)). The main idea behind this instrument is that distance to college should affect the psychic and monetary costs of attending college, but should not be correlated with labor market outcomes. The exogeneity

5. Combining these two assumptions also gives us the classic monotonicity assumption in Angrist and Imbens (1995)

of distance instruments has been assessed in Cameron and Taber (2004) and Mountjoy (2022). These studies highlight the importance of accounting for demographics, family background, and region of residence, factors we include in our analysis. In addition, we also provide suggestive evidence in section 1.4 supporting this assumption by showing that distance to college and high school test scores are uncorrelated, conditional on a set of controls for individual, high school, and neighborhood characteristics.

The exogeneity assumption does not only require that the instrument be independent from the potential earnings but also independent from the information high school students observe, conditional on the set of covariates and the potential earnings. Although this assumption is not typically discussed in the economics of education literature, it is implicitly required for the validity of the instrument. This assumption is more challenging to assess and requires us to believe that two individuals who differ only in their distance from college, but share all other observables and potential earnings, are drawing signals from the same distribution. The fact that we need the instrument to be independent from the instrument, conditional on the potential earnings, may seem reasonable if we consider that the type of information individuals observe is a function of their ability and ability is fully captured by the potential earnings of the individuals. On the other hand, this assumption may not hold if, for example, we believe that universities have outreach programs explicitly based on the distance from the college, or that information is correlated with distance from college through neighborhood components not captured in our neighborhood controls.

In addition to the requirement that the instrument satisfy the exogeneity and relevance assumptions, we also requires from our instrument to have enough variation to be able to identify the entire beliefs distribution, conditional on our set of controls. This assumption is hard to satisfy in reality and is not satisfied in our data. The variations in distance are not large enough to assure us that we can identify the entire beliefs distribution for each X. In practice, we relax this assumption and follow other papers (Carneiro et al. (2011), Carneiro and Lee (2009), Brave and Walstrum (2014), Heckman and Vytlacil (2005)) that estimate marginal treatment effects and assume that the set of covariates only operates as a mean shifter, shifting the mean of α_1 and α_0. We make this additional assumption explicit in section 1.2.5.1.

The next set of assumptions we need for identification focuses on the heterogeneity in the cost function and beliefs. As discussed at the beginning of section 1.2, conditional on the set of observables, all of our heterogeneity in choice stems from differences in the signals individuals obtain. Therefore, we assume that all heterogeneity in the cost function is observable. This assumption allows us to separate beliefs and costs in order to perform our decomposition analysis. In general, without this assumption, separating beliefs and preferences/costs from choice data is not feasible Manski (2004). To do so requires additional data directly on either beliefs, obtained from belief elicitation, or preferences, which can be measured by elaborate surveys that separate beliefs and preferences Adams-Prassl and Andrew (2019)[6]. Our decomposition analysis relies on observing both beliefs about the long-term outcomes and the realization of long-term outcomes. The requirement to observe long-term outcomes makes it difficult to obtain beliefs and realization. Therefore, we opt here to rely on the restricted heterogeneity assumption discussed here. In general, if one has access to data on beliefs they can incorporate it, as discussed in section 1.12.2.3, and allow them to better separate the cost and beliefs, at the cost of imposing another assumption.

The final key assumption our identification argument builds upon is the assumption we impose on individuals' utility and decision rule, and how they utilize the information they have. Specifically, we impose that the utility functions of individuals are strictly increasing in the expected returns and

6. We discuss briefly what additional assumptions can be used to estimate cost heterogeneity if one has access to belief elicitation data in section 1.12.2.3 in the Appendix.

that individuals have partial rational expectations, as discussed in section 1.2.2.

First, we impose that individuals have partial rational expectations to be able to measure the expectations for measured conditional expectations in the data, as we discuss below. Two things are important to keep in mind. First, imposing rational expectations implies that individuals' beliefs are not systematically detached from the underlying economy and that their beliefs correspond systematically to the underlying economy. We may want to relax this assumption, but then we would need to be precise on how beliefs are behaving in the counterfactual world where we change the underlying returns distribution, as we do in our decomposition. We also need to decide on how we measure the quality of information in this setup, and take a stance on whether this relies on incorrect beliefs or the beliefs these people would hold if they update their beliefs correctly given their available information. Therefore, relaxing the rational expectation assumptions introduces a set of decisions without a clear guideline on how to make them. We leave exploration of this to future research. Second, it's important to emphasize that for our analysis, we do not need individuals' beliefs to be correct, but we need to assume that individuals with higher expected returns also have higher beliefs. If individuals have wrong beliefs, but maintain order, such that individuals with higher beliefs do have higher returns, then then these biases are absorbed in the cost component, c, as we demonstrate below.

Next, we impose the assumption that utility of individuals is increasing in the expected returns, which allows us to pin down the belief distribution in the population by using the law of iterated expectations and the Marginal Treatment Effect curve. Specifically, consider a more general setup that allows for arbitrary utility functions. Let $u(E[Y_1 - Y_0])$ be a utility function. u can be a function of preferences, but can also include biases as discussed above. We assume that individuals opt in if $u(E[Y_1 - Y_0]) \geq 0$. If u is strictly increasing, then we have an equivalent decision rule in which $u(E[Y_1 - Y_0]) \geq 0 \iff u^{-1}(u(E[Y_1 - Y_0])) \geq u^{-1}(0) \iff E[Y_1 - Y_0] \geq c$. Now, using our instrument, we can identify the marginal treatment effect curve that maps the expected returns to each quantile of the selection variable. Specifically, denote by F the CDF of beliefs, $E[\mathcal{R}|\boldsymbol{S}]$, and by V the random variable that corresponds to the quantile of $E[\mathcal{R}|\boldsymbol{S}]$. Then we can obtain, using our instrument, the Marginal Treatment Effect curve, $E[\mathcal{R}|V=v]$. Using the law of iterated expectation, we also have:

$$\mathrm{E}\left[\mathcal{R}|V=v\right]=\mathrm{E}\left[\mathcal{R}|F^{-1}(V)=F^{-1}(v)\right]=\mathrm{E}\left[\mathcal{R}|\mathrm{E}[\mathcal{R}|s]=F^{-1}(v)\right]=F^{-1}(v)$$

This allows us to pin down the beliefs distribution F. If individuals' preferences are not strictly increasing in the returns, then the quantiles of the selection variables do not necessarily correspond to a unique set of beliefs, which prevents us from using the argument above.

In what follows we briefly go over the identification of the simpler Gaussian model, and it's important components for our analysis. Discussion on estimation is in Appendix 1.11.

1.2.5.1 Identifying the Gaussian Model Parameters

We assume we observe a set of covariates X, a continuous instrument Z and outcomes Y. Although it's not imperative for identification argument, we parametrize the cost function as a linear function of covariates

$$c(x,z)=zb_z+xb_x.$$

We assume that the distribution of $\alpha_1|X,D=1$ and $\alpha_0|D=1,X$ is observed. In the Appendix we discuss how it can be identified using panel data and additional assumptions on the wages. For our discussion α_1 and α_0 can be thought of as fixed effects, and are identified using panel data on earnings. We also assume that α_1 and α_0 are linear in covariates

$$\alpha_1 = X\beta_1 + U_1,$$
$$\alpha_0 = X\beta_0 + U_0.$$

Following our discussion on the Gaussian model, we assume that beliefs and residuals U_1 and U_0 are jointly normal, X operates only as a mean shifter and Z is independent from the potential outcomes, and $Z, X \perp\!\!\!\perp U_1, U_0$ and information $S \perp\!\!\!\perp Z, X|U_1, U_0$

$$\begin{pmatrix} U_1 \\ U_0 \\ E[\mathcal{R}|\boldsymbol{S}, x] \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \\ X(\beta_1 - \beta_0) \end{pmatrix}, \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_0 & \sigma_{1,\mathrm{E}} \\ \rho\sigma_1\sigma_0 & \sigma_0^2 & \sigma_{0,\mathrm{E}} \\ \sigma_{1,\mathrm{E}} & \sigma_{0,\mathrm{E}} & \sigma_{\mathrm{E}}^2 \end{bmatrix}\right).$$

where σ_{E} is the standard deviation of beliefs, and $\sigma_{d,\mathrm{E}}$ is the covariance between beliefs and potential earnings U_d. The decision rule is then given by

$$D = \mathbb{1}\left[E[\alpha_1 - \alpha_0 \mid \boldsymbol{S}, x] \geq c(z, x)\right] = \mathbb{1}\left[E[U_1 - U_0 \mid \boldsymbol{S}, x] \geq c(x, z) - x(\beta_1 - \beta_0)\right].$$

Using the fact that beliefs and U_1 and U_0 are jointly normal, we have that the choice probability is given by

$$P(D = 1|x, z) = \Phi\left(\frac{x(\beta_1 - \beta_0) - zb_z - xb_x}{\sigma_{\mathrm{E}}}\right). \tag{1.4}$$

Notice that in general this is not enough to identify the cost function of parameters, as all parameters are identified up to scale. In addition, covariates can play a dual role, both affecting the outcome variable and controlling the cost. Therefore, we need to identify the scale parameter and the coefficients β_1 and β_0. To identify β_1 we use the standard Heckman Correction argument for Gaussian selection model (Heckman (1979)). Specifically, using the fact that U_1, U_0 and beliefs are jointly Gaussian, we have that

$$E[\alpha_1|D=1,X] = E[\alpha_1+U_1] = E[\alpha_1+U_1] = X\beta_1 + E[U_1|D=1,X],$$

where $E[U_1|D=1,X] = \frac{\sigma_{1,\mathrm{E}}}{\sigma_{\mathrm{E}}} \frac{\phi\left(\Phi^{-1}(1-P(D=1|x,z))\right)}{1-\Phi(\Phi^{-1}(1-P(D=1|x,z)))}$. We can follow the same argument to identify β_0, and using the fact that $E[\alpha_0|D=0,X] = \frac{\sigma_{0,\mathrm{E}}}{\sigma_{\mathrm{E}}} \times -\frac{\phi\left(\Phi^{-1}(1-P(D=1|x,z))\right)}{\Phi(\Phi^{-1}(1-P(D=1|x,z)))}$. Denote the coefficient of the inverse mills ratio as $\gamma_1 = \frac{\sigma_{1,\mathrm{E}}}{\sigma_{\mathrm{E}}}$ and $\gamma_0 = \frac{\sigma_{0,\mathrm{E}}}{\sigma_E}$, and notice that we can identify σ_1 and σ_0 using the joint distribution of choice and earnings

$$f(D=1,\alpha_1,z,x) = \left(1-\Phi\left(\frac{\Phi^{-1}\left(1-P(D=1|x,z)\right) - \frac{\gamma_1}{\sigma_1^2}\left(\frac{\alpha_1-x\beta_1}{\sigma_1}\right)}{\sqrt{(1-(\frac{\gamma_1}{\sigma_1})^2))}}\right)\right)\phi\left(\frac{\alpha_1-x\beta_1}{\sigma_1}\right)\frac{1}{\sigma_1}. \tag{1.5}$$

and similarly for σ_0. Finally, in order to get σ_{E}, we can use two facts. First, notice that that the covariance of beliefs and returns equal to the variance of returns, $\mathrm{Cov}(U_1-U_0, \mathrm{E}[U_1-U_0|\boldsymbol{S},x]) = \mathrm{Var}(\mathrm{E}[U_1-U_0|\boldsymbol{S},x])$. To see that notice that we can decompose returns as

$$U_1 - U_0 = E[U_1-U_0|\boldsymbol{S},x] + r,$$

where r is the residual from projecting U_1-U_0 on $\boldsymbol{S},X$, and satisfies $\mathrm{Cov}(E[\mathcal{R}|\boldsymbol{S},x],r)=0$. Second, we have that $\mathrm{Cov}(U_1-U_0, E[\mathcal{R}|\boldsymbol{S},x]) = \mathrm{Cov}(U_1, E[\mathcal{R}|\boldsymbol{S},x]) - \mathrm{Cov}(U_0, E[\mathcal{R}|\boldsymbol{S},x]) = \sigma_{1,\mathrm{E}} - \sigma_{0,\mathrm{E}}$. Combining these two facts we have from the two coefficients on the control function in the potential earnings regression

$$\gamma_1 - \gamma_0 = \frac{\sigma_{1,\mathrm{E}} - \sigma_{0,\mathrm{E}}}{\sigma_{\mathrm{E}}} = \frac{\sigma_{\mathrm{E}}^2}{\sigma_{\mathrm{E}}} = \sigma_{\mathrm{E}}.$$

which concludes the identification argument for all the component we need for our decomposition.

1.2.6 What can be learned on the prior beliefs on the correlation ρ

The data we use in our empirical application and the restrictions the choice model outlined above implies, do not allow us to identify the prior distribution individuals have over the correlation parameter ρ. In this section, we discuss what *can* be learned on the prior beliefs given our model and data. Our main results here shows that given our identification results above, we can bound the set of feasible ρs, from the individuals' perspective. To show this, we first start by showing how the variance of beliefs introduces some restrictions on the set of feasible ρs. We then continue to show how we can identify the set of feasible ρs under an additional assumption on the quality of information individuals have on the marginal distribution. Finally, we demonstrate how we can bound on the quality of information parameter in equation 1.2.

1.2.6.1 Restrictions on the Correlation Parameter

Our theoretical framework implies some constraints on the correlation between U_1 and U_0, that is informed by our model that implies some selection on returns. First, as it well known, the variance of beliefs about returns is bounded from above by the actual variance of returns (e.g Gentzkow and Kamenica (2016)), which implies that the following inequality must hold:

$$\mathrm{Var}(E[\mathcal{R}|\boldsymbol{s}] \leq \sigma_1^2 + \sigma_0^2 - 2\rho\sigma_1\sigma_0.$$

This restriction is a generalization of the known fact in the standard Roy model (Roy (1951)) with complete outcome information, where the joint distribution of potential outcomes is point-identified (Heckman and Robb (1985)). If we assume agents have complete information, the inequality holds with equality and we can identify the joint distribution of potential earnings. If we maintain that agents select based on outcomes but have incomplete information, we can use the above inequality to bound the correlation between potential outcomes.

We can further restrict the bounds using the fact that we can identify the covariance between beliefs, $\mathrm{E}[\alpha_1 - \alpha_0 | s, x]$ and U_1 and U_0. To do so we use the fact that the covariance matrix must remain positive semi-definite, we therefore restrict the set of possible ρ to values that keep the following covariance matrix positive semi-definite,

$$\mathrm{Cov}(\boldsymbol{\alpha}, \mathbb{E}) = \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_0 & \sigma_{1,\mathbb{E}} \\ \rho\sigma_1\sigma_0 & \sigma_0^2 & \sigma_{0,\mathbb{E}} \\ \sigma_{1,\mathbb{E}} & \sigma_{0,\mathbb{E}} & \sigma_{\mathbb{E}}^2 \end{bmatrix}.$$

1.2.6.2 Identifying the Set of Feasible ρs from the High School Students' Perspective

Our measure of information quality depends on the set of feasible values of ρ taken from the perspective of high school graduates. One way to identify this set is to assume that the set of feasible correlations obtained using our bounding method above is the same as the set of feasible ρs from the high school students' perspective. This would be the case if, for example, high school graduates observe only a scalar signal, such that their beliefs are an injective function of their signal. This assumption might be very restrictive, and in general, individuals are likely to have access to various sources of information and observe multiple signals. Without additional assumptions, our model is not restrictive enough to pin down the set of feasible ρs a high school student may consider, as the correlation between the signals and potential earnings can induce additional restrictions on ρ that are not captured by the argument above. To overcome this, we first show that with an additional assumption on the quality of information individuals have on the marginal U_1, we can identify the set of feasible ρs from the high school graduates' perspective.

We start by defining what is the set of feasible ρs from the high school student perspective. We then define the set of feasible set of ρs from the the econometrician's perspective, under an additional assumption on the quality of information on the marginals. Let $\boldsymbol{S}$ be a vector of signals high school students have. The set of feasible ρs is the set of ρs that keep the covariance matrix of $\boldsymbol{S}, U_1, U_0$ positive semi-definite (PSD). We say that a ρ is feasible from the perspective of high school graduates if the covariance matrix between $\boldsymbol{S}, U_1, U_0$, where the correlation between U_1 and U_0 is ρ, is also PSD.

Next, we consider what may be potentially known to the econometrician. Denote by R_1^2 the quality of information high school students have on U_1, i.e., $R_1^2 = \frac{\text{Var}(E[U_1|S])}{\text{Var}(U_1)}$. The next lemma shows that the covariance matrix between $U_1, U_0, E[U_1|s]$, and $E[U_0|s]$ is identified, up to ρ, for a given R_1^2.

Assume R_1^2 is known, then all components of the covariance matrix between U_1, U_0, $\mathbb{E}[U_1|\boldsymbol{S}]$, and $\mathbb{E}[U_0|\boldsymbol{S}]$ are identified up to ρ. The proof is in Appendix 1.10.1 and builds on the identification results of the model, as discussed above, and the fact that the conditional expectation $\mathbb{E}[U_d|\boldsymbol{S}]$ in the Gaussian model is linear. Finally, we then say that ρ is feasible from the econometrician's perspective, for a given value of R_1^2 the implied covariance matrix between $U_1, U_0, E[U_1|s]$, and $E[U_0|s]$ is PSD.

The following proposition shows that if ρ is feasible from the econometrician's perspective, with the assumption on R_d^2, then it is also feasible from the high school graduates' perspective.

Proposition 1. Fix R_1^2. A ρ is feasible from the high school graduate perspective if and only if it is feasible from the econometrician's perspective.

The proof is in section 1.10.1 in the Appendix, and builds on the linearity and of the normal conditional expectation and properties of PSD matrices. Proposition 1 demonstrates that, given an assumption on the quality of information individuals have on one of the marginals, we can identify the set of feasible ρs from the high school graduate perspective. Furthermore, we also note that the set of feasible ρ is a closed interval.[7] Therefore, we can describe the set by its boundaries ρ_{min} and ρ_{max}. Finally, Proposition 1 shows that we can identify the set of feasible ρs, from the perspective of high schools students, under an assumption on the quality of information they hold on the U_1, R_1^2. Since we do not know the R_1^2, in the results section, when we use the support to inform our choice of prior of ρ, we construct bounds on the information quality, R^2, from equation 1.2, by exploring all values of $R_1^2 \in [0, 1]$.

7. To see that, remember that covariance matrix

$$C = \begin{pmatrix} \Sigma_S & \Sigma_{S,1} & \Sigma_{S,0} \\ \Sigma_{S,1}^T & \sigma_1^2 & \rho\sigma_1\sigma_0 \\ \Sigma_{S,0}^T & \rho\sigma_1\sigma_0 & \sigma_0^2 \end{pmatrix}$$

is positive semi-definite (PSD) if and only if Σ_S is PSD and the Schur complement of Σ_S in C is also PSD. Σ_S is PSD by construction. The Schur complement, denoted as SC, is given by

$$SC = \begin{pmatrix} \sigma_1^2 & \rho\sigma_1\sigma_0 \\ \rho\sigma_1\sigma_0 & \sigma_0^2 \end{pmatrix} - \begin{bmatrix} \Sigma_{S,1}^T & \Sigma_{S,0}^T \end{bmatrix} \Sigma_S^{-1} \begin{bmatrix} \Sigma_{S,1} \\ \Sigma_{S,0} \end{bmatrix}.$$

The SC is PSD, if $u^T SCu \geq 0$ for any vector u. We can demonstrate that this holds if $k_2x^2 + (k_1 - \rho)x + k_0 \geq 0$, where k_0, k_1, and k_2 are constants determined by the SC elements and $x = \frac{u_1}{u_0}$. To ensure that this expression is always positive, we can use the quadratic formula and require that $(k_1 - \rho)^2 - 4k_2k_0 \leq 0$. This expression is a convex parabola in ρ, that intersects with the constant $4k_2k_0$ at at most two intersection points. Any ρ between these two intersection points satisfies the requirement and maintain that matrix C is PSD. Consequently, it is sufficient to describe the set of feasible ρ values by these two boundary points. If the parabola does not intersect with $4k_2k_1$ then there is no ρ that keeps matrix C a PSD matrix.

1.3 Data

Our empirical application investigates the factors contributing to the college attendance gap between Hispanic and White students. We concentrate on Texas, where there are large and comparable Hispanic and White populations, but they differ substantially in their choices. Utilizing the methods described in Sections 1.2.4, we decompose the attendance choices and assess the influence of informational differences. We start by describing the data and then discuss the model results. The following section describes the data and variables we use throughout our analysis

1.3.1 Data Sources and Sample Construction

Our empirical study leverages a series of confidential administrative databases from the state of Texas, the second most populous in the U.S. with a sophisticated higher education system that engages a substantial portion of its populace, including over one million high school students (Agency (2023)). Additionally, Texas have a significant Hispanic demographic, comprising around 12 million individuals in 2022, or about 40% of the state's total population, matched by a 40% representation of White population.

The study combines data from several Texas agencies. The primary dataset is procured from the Texas Education Agency (TEA), offering demographic details of all Texan high school students. This dataset is enriched with school characteristics from the National Center for Education Statistics (NCES), which provides a broader picture of Texas high schools. We incorporate assessments from the Texas standardized testing program, which evaluates public primary and secondary school students' competencies in various grades and subjects. Further, we integrate data concerning college enrollment decisions from the Texas Higher Education Coordinating Board (THECB), supplemented by information from the Integrated Postsecondary Education Data System (IPEDS). Finally, the Texas Workforce Commission (TWC) supplies data on post-high school earnings, completing our comprehensive dataset.

In constructing our control variables, we follow the approach used by Mountjoy (2022), utilizing three types of covariates: student-level demographics, school characteristics, and neighborhood characteristics. For student-level demographics, we include categorical variables for gender, eligibility for free or reduced price lunch as a proxy for economic disadvantage, and an indicator for graduation under one of three programs: the Distinguished Achievement Program, Recommended High School Program, or the Minimum High School Program, which reflect the various graduation tracks in Texas. In some of our analyses, we use test scores from Texas Assessment of Knowledge and Skills (TAKS) tests. We consider test scores from the exit exams in English-Language-Arts (ELA), which capture language skills, and Math test scores, these tests were held consistently across our three cohorts of interest. We then create a single measure of test scores by combining them in a one-factor model separately by cohort and normalize this factor to within-cohort percentiles. These high-stakes tests, which imply that they are likely to be indicative of student ability .Passing these exit-level test is a graduation prerequisite for Texas high school seniors in their junior and senior years.

For high school-level controls, we utilize NCES Common Core data, which incorporates the geographic locale code. This code categorizes urbanization into twelve detailed categories using Census geospatial data. Additionally, we include the distance to two-year colleges and an indicator denoting whether the school is classified as a Vocational Education School. Vocational schools are identified as those that provide formal training for semi-skilled, skilled, technical, or professional occupations to students of high school age who may opt to enhance their employment prospects, possibly instead of preparing for college admission. Controls also account for the local influence of the oil and gas industry, by measuring the long-term share of oil and gas employment at the high school level, employing NAICS industry codes from TWC workforce data. We normalize this measure of oil and gas employment by ranking it and control for its effects using a third-degree polynomial in our analysis of school characteristics.

Neighborhood characteristics include the 62 Texas commuting zones using the year-2000 mapping provided by the U.S. Department of Agriculture's Economic Research Service. We also construct an index of neighborhood quality, akin to the test score measure: We combine the tract-level Census measures of median household income and the percentage of households below the poverty line with the high school-level percentage eligible for free/reduced-price lunch into a one-factor model, then normalize this neighborhood factor to the within-cohort percentile. When controlling for neighborhood characteristics in the following discussion, we control for the third-degree polynomial of the neighborhood factor.

As outlined in section 1.10.2 in the Appendix, nonparametric identification necessitates an instrument. We employ the measure of proximity to the nearest 4-year colleges, calculating ellipsoidal distances between the coordinates of all Texas public high schools (sourced from NCES CCD) and those of all Texas postsecondary institutions (from IPEDS). We determine the minimum distances within 4-year sectors for each high school. To supplement some missing distances, we refer to Mountjoy (2022), which involved manual collection of location data by verifying each college's institutional profile. We adopt the same methodology for the variable of distance to 2-year colleges.

We limit our sample to cohorts from 2003 to 2005 to ensure a long time horizon. This approach, leveraging our earnings data, allows us to observe outcomes 16 (for the cohort of 2003 and 2004) and 15 years (for the cohort of 2005) into the future, thus better understanding the incentives faced by these students. Additionally, the Texas Higher Education Coordinating Board (THECB) has provided data on students attending four-year colleges, including both private and public institutions, starting from 2003. We further narrow our sample to high school students who are not enrolled in special education programs, are between the ages of 17 and 18 in the 12th grade, and have graduated from high school with at least the minimum requirements. As with any study focused on a specific state, there is a risk of out-migration; however, Texas has one of the lowest out-migration rates in the U.S. (Times (2014)). Following Mountjoy (2022), we also limit our test factor to individuals with grades below the 80s percentile. As Mountjoy (2022) discusses, high school students with a test score factor higher than the 80th percentile are more likely to enroll in out-of-state colleges. Figure 1.10 in the Appendix further illustrates that these individuals are more likely to have missing earnings data.

1.4 Summary Statistics and Empirical Patterns

Table 1.5 in the Appendix presents summary statistics for the analysis cohorts. The table shows substantial disparity in socio-economic backgrounds among the groups. A significant proportion of Hispanics originate from low-income families, necessitating reduced-price or free meals. They also live in census tracts with higher unemployment rates and a greater proportion of families below the poverty line. Over 58% of Hispanics attend Title I schools, markedly more than their White counterparts. Conversely, regarding the programs offered at these schools, there is no substantial difference in the distribution. Similarly, there is no significant difference in how schools inform students about the oil industry; the proportion of high school graduates working in the oil and gas industries over the long term is similar. Geographically, Hispanics are more likely to reside in urban areas, while Whites predominantly live in suburban and rural areas. Furthermore, in terms of proximity to colleges, Hispanics tend to live nearer to both four-year and two-year colleges compared to non-Hispanic Whites.

In what follows, we delve deeper to describe the college attendance gap and the two driving mechanisms: earnings and information.

1.4.1 College Attendance

The first row of Table 1.5 in the Appendix shows that the choice gap in the decision to attend a four-year college in the first year after high school graduation between Hispanics and Whites is 9%. Table 1.11 in the Appendix examines the extent to which observable factors contribute to this disparity. The first row adds control for individual characteristics. Controlling for neighborhood characteristics increases the average choice gap to 13%, implying that the choice gap between Hispanics and Whites who reside in similar neighborhoods is larger than the average choice gap in the population. Controlling for individual characteristics reduces the remaining gap back to 8.6%, controlling for school characteristics does not change the gap by much and reduces it to around 7.6%. Finally, controlling also for test scores reduces the gap to 4.28%, implying that test scores help explain a large portion of the choice gap.

Figures 1.2 and 1.3 illustrate that there is high dispersion in both Whites' and Hispanics' likelihood of attending college. These figures plot histograms of the propensity scores for Hispanics and Whites attending college, estimated using a Probit model with our control set and the distance to a four-year college. Firstly, they reveal a large overlap in propensity scores, as required for our identification argument, as discussed in Section 1.2.5 and Appendix 1.10.2. Furthermore, the figures demonstrate that Whites are more likely to attend college, ex-ante, based on their characteristics, as for both college-goers and non-college goers, the distribution of propensity scores for Whites is more skewed to the right.

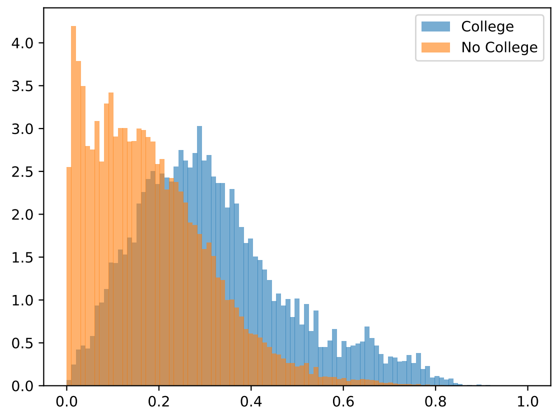

Figure 1.2: Propensity Scores - Hispanics

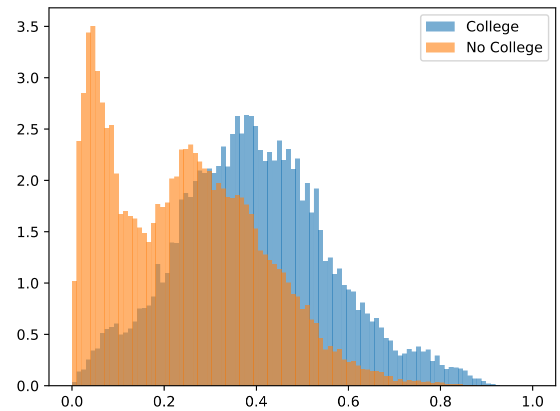

Figure 1.3: Propensity Scores - Whites

1.4.2 Earnings

We now turn to focus on the differences in earnings distributions between Hispanics and Whites. Tables 1.6 and 1.7 in the Appendix show the average quarterly earnings for Whites and Hispanics at various intervals post-graduation. Generally, wages are on an upward trend over time, albeit at a decreasing rate. The tables show that mean earnings of Whites are higher than those of Hispanics for any period after high school graduation. When splitting the earnings by college goers and non-college goers, we can see that after 15 years, Whites' earnings for both college goers and non-college goers are higher. The table also shows that in the first years after high school graduation, the differences between Hispanics' and Whites' earnings are small.

For our discussion, the most important component is the differences in earnings between college attenders and non-college attenders, across Hispanics and Whites. Figure 1.4 explores this difference. The figure plots the coefficient for attending a four-year college for both Hispanics and Whites, controlling for cohort fixed effects. The figure shows that the gap in earnings between first-year college goers and non-college goers increases over time for both Hispanics and Whites. Notably, this difference widens in the first five years post-graduation and then stabilizes at around \$500, which is approximately 6% of the average quarterly earnings for Hispanics 14-16 years after graduation. Figure 1.5 introduces our set of individual, school level, and neighborhood level controls. The figure shows that adding these controls reduces the levels but does not affect the gap, demonstrating that the gap in earnings is not fully explained by these controls.

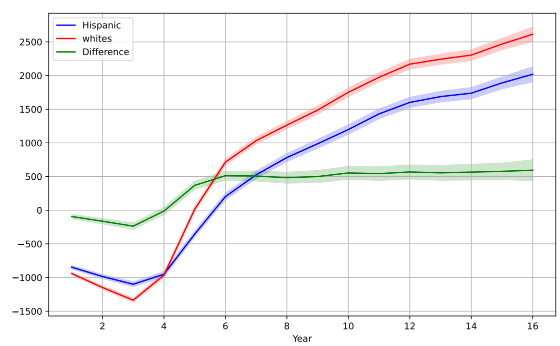

Figure 1.4: Raw difference in Mean Wages, w/out controls

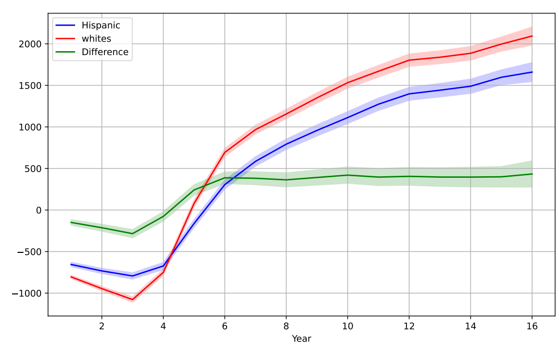

Figure 1.5: Raw difference in Mean Wages, with controls

Note: Figure 1.4 plots the coefficient for attending a four-year college for both Hispanics and Whites, controlling only for cohort fixed effect. The Coefficient for 16 years after college is using only two cohort, 2003-2004. Figure 1.5 plots the same coefficient, with all the added controls, as discussed in section 1.3

Within the framework of our model, these differences suggest that Hispanic high school graduates may have less incentive to attend college compared to their White counterparts. However, this observed gap could be attributable to selection bias rather than reflecting the actual returns faced by the high school graduates. To overcome this selection effect, we utilize the distance to college from high school as an instrument in a Two-Stage Least Squares (TSLS) analysis.[8]

8. The use of distance to college as an instrumental variable has been prevalent in the literature that estimates returns to education. See Card (1995) and its subsequent application in works such as Carneiro et al. (2011), Carneiro and Lee (2009), Kapor (2020), Abdulkadiroğlu et al. (2020), Mountjoy (2022).

First, to examine the instrument's validity, we explore its relation with test scores. As we discuss below, test scores are both associated with potential outcomes and with the decision to go to college. Therefore, if the exclusion restriction holds, we do not expect that distance to college should be correlated with test scores, conditioned on our set of controls. Table 1.8 in the Appendix examines the correlation between the instrument and test scores. Initially, without our set of controls, test scores show a significant correlation with the instrument. After including individual characteristics, this correlation persists, which might indicate that spatial sorting is non-random and likely tied to other factors that influence both outcomes and information. Subsequent rows in the table introduce more controls for school and neighborhood characteristics, which largely account for the initial correlation, rendering the coefficient on distance nearly null, indicating that the instrument may be valid.

We next examine the relevance assumption needed for the instrumental variable. Table 1.9 in the Appendix shows a strong first stage: the influence of distance to college on the likelihood of attending a four-year college immediately after graduation. Controlling for our set of controls, we see that an increase of one mile in distance to college decreases the likelihood of college attendance by 0.2% for Hispanics and 0.1% for Whites. The magnitude of this effect remains relatively stable upon the inclusion of different controls. Keep our discussion in section 1.2.3 in mind. The fact that the effect of the distance to college is stronger for Hispanics indicates that either their beliefs are more dispersed, which implies that they have better information quality, or that the difference $\mu_{\mathcal{R}} - C$ is higher.

Table 1.1 presents the results from the TSLS regression that instruments the treatment effect using the distance to college instrument and includes all controls. It shows that, after adjusting for selection, the average effect for Hispanics is negligible, persisting up to 16 years post-high school graduation. For Whites, on the other hand, there is a gradual effect that mirrors the earnings dynamics depicted in Figure 1.5. These findings suggest that the returns for Hispanics are generally much lower, potentially diminishing the incentive to pursue higher education.

	All	Hispanics	Whites
Avg. Wage 8-10	245.0	707.0	-1108.62
	(1194.0)	(1237.0)	(2028.0)
	245206	103198	142008
Avg. Wage 10-12	875.0	305.0	521.0
	(1436.0)	(1468.0)	(2295.0)
	239307	101284	138023
Avg. Wage 12-14	1552.0	255.0	2380.0
	(1531.0)	(1550.0)	(2370.0)
	233091	99428	133663
Avg. Wage 14-16	2605.0	377.0	5156.0
	(1632.0)	(1745.0)	(2424.0)
	149498	63271	86227

Table 1.1: Returns - Two Least Squares

Note: This table presents the results from a Two-Stage Least Squares (TSLS) regression of college attendance on earnings. Earnings are measured in periods of 8-10, 10-12, 12-14, and 14-16 years after the students' high school graduation. We instrument college attendance using the distance to the nearest college and control for individual, school, and neighborhood characteristics, as discussed in Section 1.3. For the 8-14 year period post-graduation, we include cohorts from 2003-2005. For the 14-16 year period, we include only the 2003-2004 cohorts due to data limitations.

Finally, our measures of quality focus on the amount of variance in earnings that information can explain. Table 1.6 demonstrates that Whites have higher variability in earnings compared to Hispanics at each point in time after graduating, hinting that more information is needed to better predict Whites' earnings than Hispanics'. We explore this notion further in Figure 1.13 in the Appendix, where we demonstrate that not only are wages more variable, but there is also higher variability in the industries in which Whites work. The figure plots Shannon's entropy for the 2-digit NAICS industry codes in which Hispanics and Whites are employed each year after high school graduation. The figure shows that, throughout their lives, Whites are less concentrated in specific industries compared to Hispanics. This also supports that it's harder to predict Whites' later-life outcomes compared to Hispanics.

1.4.3 Information

To get a sense of the quality of information is challenging, as we do not observe in the data the pieces of information high school students have access to. We therefore consider specific signals we can observe in our data, or in auxiliary data sets. Specifically, we first examine how informative the information contained in school performance measures is. This is information we can observe in our data, and students are likely to hold and use when making decisions on whether or not to attend college. We then continue to describe survey results that demonstrate that Hispanic and White students utilize similar sources of information in their decisions related to career and education choices.

Test scores and school performance provide important information for high school students for their decision-making process. Grades act as sources of information and signals available to students before making a decision. From this perspective, agents receive grades and use them to form projections about the utility of these grades. Consequently, we also examine whether grades convey informative signals about returns and whether there exists disparities in quality between Whites and Hispanics.

Table 1.5 reveals a notable gap in academic readiness between Hispanics and Whites, as evidenced by exit exam grades. To what extent does this gap contribute to the overall disparity? We first show that grades and test scores are likely to affect choices, as discussed above. The final row in Table 1.11 in the Appendix demonstrates that when we account for our measure of test scores, the gap narrows to 4.8%, implying that at least some of the gap is driven by differences in test scores. As the decision to attend college is made after test scores are known, this suggests that test scores themselves are used in the decision process. Furthermore, Table 1.15 demonstrates that grades are significant in explaining choices. In a Probit model predicting these choices, the inclusion of grades increases the Area Under the Curve (AUC) from 0.74 to 0.77 for Whites and from 0.75 to 0.8 for Hispanics. This magnitude of increase is comparable to that observed when adding school and neighborhood characteristics to individual characteristics, rising from 0.68 to 0.75 for Hispanics and from 0.67 to 0.74 for Whites. These findings again imply that high school graduates likely consider exit exam grades and their informational value in their college enrollment decisions.

Are grades informative on returns? We first explore whether grades are likely to contain information about returns. To ascertain whether grades predict earnings and returns, Figure 1.11 in the Appendix illustrates the relationship between earnings and grades for both college attendees and non-attendees. The figure shows that for both Hispanics and Whites, higher grades correlate with increased earnings, irrespective of college attendance. Additionally, as grades increase, the earnings gap widens between those who attend college in their first year and those who do not. This is supported by the regression in Table 1.12 in the Appendix, which reveals that a one-unit increase in test scores raises the raw gap by approximately \$16, controlling for our set of controls. Both figures and the regression table suggest that the difference in informativeness of test scores across the two group is relatively small.

The relationship between school informativeness is further examined in Table 1.2, which we discuss further in Section 1.5. This table presents the out-of-sample R^2 from a model that employs Extreme Gradient Boosting to predict earnings based on students' course-taking patterns and the pass-fail indicator for Hispanics and Whites. The R^2 values are remarkably similar for both groups. This implies that the quality of information from school performance measures is comparable for Whites and Hispanics.

Finally, to explore what other sources of information are used by high school students, we use a

			In Sample R^2	Out of Sample R^2
Fixed Effects	No College	All	0.19	0.11
		Hispanic	0.17	0.09
		Whites	0.18	0.10
	College	All	0.15	0.09
		Hispanic	0.14	0.09
		Whites	0.11	0.06
W/O Fixed Effects	No College	All	0.20	0.10
		Hispanic	0.18	0.09
		Whites	0.20	0.09
	College	All	0.15	0.10
		Hispanic	0.16	0.09
		Whites	0.13	0.08

Table 1.2: School Informativeness - R^2

Note: This table displays the in-sample and out-of-sample R^2 values for a model predicting average earnings 12-14 years post high school graduation. The No-FE rows ("No Fixed Effect") incorporates individual characteristics (as detailed in Section 1.3), test scores from exit exams in math and English comprehension, and indicators for each course taken during the three years of high school, including pass/fail status, taken from the Texas Education Agency data. The FE rows ("Fixed Effect") additionally includes a high school indicator variables, controlling for the impact of different high schools. Estimation is conducted using XG-Boost, with parameter selection via Parallelizable Bayesian Optimization, as implemented in the R package "Parallelizable Bayesian Optimization."

survey conducted by the Texas Higher Education Opportunity Project[9]. Table 1.13 in the Appendix shows that Hispanic high school students are slightly more likely than their White peers to approach and discuss with the school counselor about education and career decisions. Specifically, 56% of Hispanics discuss their school counselor about career options vs. only 45% of Whites. Similarly, 61% of Hispanics discuss with their school counselor about college options, vs. 58% of Whites. Furthermore, Table 1.14 shows that the number of yearly interactions with the school counselor on these and other matters is almost the same across both Hispanics and Whites, indicating that the nature of interaction across the two groups is similar. Table 1.15 in the Appendix shows that Hispanics are slightly more likely to seek advice from their parents about educational and career

9. A more detailed description of Texas Higher Education Opportunity Project can be found in 1.13 in the Appendix.

decisions. These indicators together demonstrate that Hispanics and Whites turn to the same type of information sources for information.

These results indicate that Hispanics and Whites encounter varying distributions of returns. However, the quality of information available to them through the school system does not significantly differ. This motivates the utilization of our model to gain a deeper understanding of how these differences contribute to the choice gap.

1.5 Model Results

In this section, we estimate the model outlined in section 1.2.2 and discuss the implications of the estimated parameter for the role of information in determining the gap. Our analysis assumes that individuals are primarily concerned with their quarterly earnings 12-15 years post-graduation. As demonstrated in table 1.1, positive returns to college education starts approximately after 12 years. Consequently, we average the quarterly earnings within this 12-15 year period. This approach enables us to use data from our three cohorts and effectively capture the structural components, averaging over a long period. Detailed discussion on the estimation method is in Appendix 1.11.

We start our analysis by examining the relationship between the perceived cost of attending college and beliefs among Hispanic and White students. Figures 1.7 and 1.6 present histograms of the estimated costs for these groups, revealing that Hispanic students generally face lower attendance costs. As discussed in section 1.2, these costs encompass barriers to entry, such as credit constraints or discrimination, and are also influenced by preferences shaped by social norms and other factors. Table 1.3 further shows that the average cost for Hispanic students corresponds to \$1,199 of their quarterly earnings, compared to \$2,879 for White students. In addition the to cost, figures 1.6 and 1.7 also explore the distribution of conditional returns $\mathrm{E}[\alpha_1 - \alpha_0 | x]$, which represent the mean beliefs about returns for individuals with characteristic x. The two figures demonstrate that White students exhibit significantly higher expected returns than Hispanic students.

Table 1.3 further complements this analysis, showing that the average beliefs on returns are lower than the actual average return for both groups. Specifically, the gap between the mean costs and mean beliefs about returns is narrower for White students (\$949) than for Hispanic students (\$2256). This implies that to achieve parity in choice across groups, the information quality of Hispanics have should be much better than those of Whites, as discussed in section 1.2.

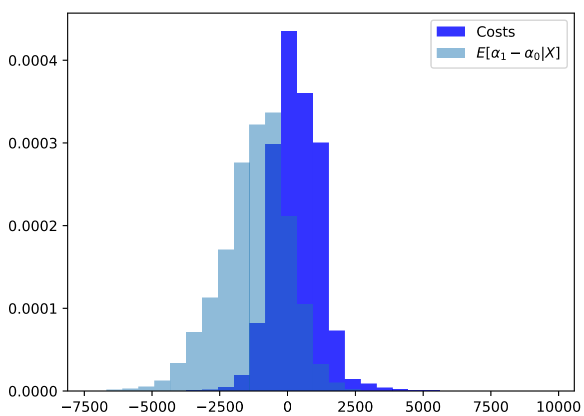

Figure 1.6: Hispanics Costs and Beliefs

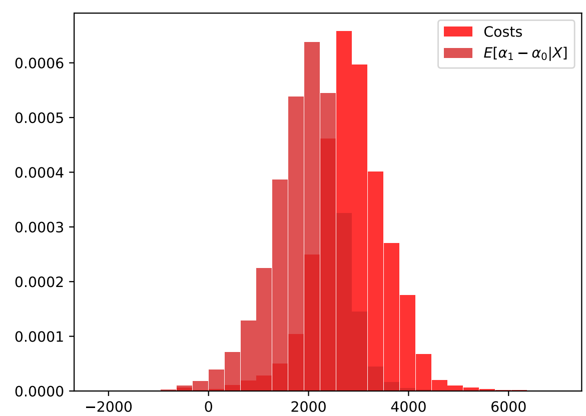

Figure 1.7: Whites Costs and Beliefs

Note: Figures 1.7 and 1.6 present histograms of the estimated costs and conditional priors $E[\alpha_1 - \alpha_0|X]$ for Hispanics and Whites, respectively. These estimates are derived according to the model discussed in Section 1.2. The parameters α_1 and α_0 represent the average quarterly earnings of high school students 12-15 years after graduation.

We now turn to look at the estimated distributions of beliefs on $U_1 - U_0$. Table 1.3 presents the estimated variance of α_1, α_0 residuals, and beliefs for the two groups. The variance in beliefs among Hispanics is notably higher than that among Whites, and the variances for the residuals of α_1 and α_0 the means are generally higher among Whites, although we can't rule out statistically that they are equal. These differences suggest that the quality of information, as measured by R^2, is the same or lower for Whites than it is for Hispanics. If the residual variance of returns for Whites is higher or the same as that for Hispanics, this implies that choice outcomes are less predictable for Whites. In both cases, the quality of information on returns hinges on the covariance structure of U_1 and U_0. Figure 1.14(a) in the Appendix shows plots the estimated CDF of the beliefs distribution, conditioned on the average covariates, and figure 1.14(b) shows the CDF where we fix all covariates and constant to zero. The figure shows that for both Hispanics and Whites, the beliefs are systematically higher for the average White high school student. Concentrating on the CDF's shape when $X = 0$, we can again see that for White and Hispanics Students with the same observables, the beliefs of Whites are less dispersed.

	P(D=1)	σ_{E}	σ_0	σ_1	Avg. Cost	$E[\alpha_1]$	$E[\alpha_0]$	$E[\alpha_1 - \alpha_0]$
Hispanics	0.21	2381	4490	6264	1199	6658.0	7715	-1057
		(657.0)	(125.0)	(818.0)	[889]	(1795.0)	(1843.0)	[2573.0]
Whites	0.29	1414	5577	6316	2879	10871	8942.0	1930.0
		(873.0)	(155.0)	(491.0)	[693]	(2149.0)	(2211.0)	[3083.0]

Table 1.3: Model Parameters

Note: The table displays model parameters estimated using average quarterly earnings 12-15 years after high school graduation. Standard errors for these parameters are presented in round parentheses ( ). Standard deviations of the costs and beliefs are indicated in square brackets [ ].

1.5.1 Measuring the Contribution of Information Differences to the Choice Gap

In this section we explore the decomposition results. We first consider the analysis of our results for the case in which $\mu_\rho = 0$, we then explore how robust our results, under different assumptions on the the prior distribution.

1.5.1.1 Main Results

Our objective is to explain the almost 8% gap in college attendance decisions between Hispanics and Whites. Table 1.4 explores the decomposition exercise, focusing on how much of the gap is explained by differences in information quality between the two groups. Row 1 of the table shows the information and returns channel, under our main assumption that $\mu_\rho = 0$. The table shows that most of the choice gap is driven by the returns channel. Specifically, the information channel constitutes -97% of the choice gap, which implies that eliminating the differences in information would result in increasing the gap by 7.7 percentage points. On the other hand, the returns channel consittute 197% of the gap, and eliminating it reduces reduces the choice gap by 15.6 percentage points. Therefore the finding shows that fully eliminating the differences in returns would eliminate the choice gap and reverse it, causing the Hispanic share to surpass the share of Whites who choose to enroll.

The decomposition shows that most of the choice gap is driven by differences in returns, with information playing a small role in shaping the gap. Moreover, current differences in information help to mitigate the gap, and if both groups had the same quality of information, the gap is likely to double in size. Therefore, the results indicate that a policy aimed at reducing differences in information quality is likely to be less effective in reducing the gap than a policy that targets differences in the returns to college for the two groups.

1.5.1.2 Robustness Analysis

Our decomposition results are partially based on the assumption we have on the prior beliefs on ρ. In this section, we explore the robustness of our results concerning this assumption. To do so, we consider two approaches. First, we relax the assumption that we know the support of feasible

	Information Channel	**Returns Channel**
1) $\mu_\rho = 0$	-0.077 (-97.116%)	0.156 (197.0%)
2) All Possible R_1^2		
LB, CF= 0.41	-0.121 (-152.532%)	0.2 (253.0%)
UB,CF= 0.243	0.046 (57.807%)	0.033 (42.193%)
3) $R_1^2 \leq 0.3$		
LB, CF= 0.369	-0.079 (-100.15%)	0.159 (200.0%)
UB,CF= 0.342	-0.053 (-66.8%)	0.132 (167.0%)
4) $R_1^2 \leq 0.5$		
LB, CF= 0.378	-0.089 (-112.35%)	0.169 (212.0%)
UB,CF= 0.325	-0.036 (-45.38%)	0.115 (145.0%)
5) Unrestricted Mean Beliefs		
LB, CF= 0.456	-0.167 (-209.918%)	0.246 (310.0%)
UB,CF= 0.189	0.1 (126.0%)	-0.021 (-25.952%)

Table 1.4: Main Decomposition

Note: This table shows the main decomposition results. Row 1 shows our main results, the decomposition of the choice gap into the information channel and returns channel. Rows 2-4 show the upper bound (UB) and lower bound (LB) of the information channel under the different assumptions on the quality of information individuals have on their college earnings (R_1^2), as discussed in the main text. The bounds on the Counterfactual (CF) share of Whites who would go to college if they had the information quality of Hispanics are shown for each case. Row 5 shows the lower and upper bounds of the information channel and returns channel in the case where we only restrict the mean prior beliefs to lie on the feasible set ρ.

ρ's but maintain the assumption that the beliefs are centered in the middle of the support, from the individual perspective, as discussed in section 1.2. Next, we relax the assumption that mean priors are at the middle of the support and explore how this may affect the returns and information channels.

We start by imposing our insight on what can be learned about the support of ρ from the perspective of the individuals. Row 2 in Table 1.4 shows the upper and lower bounds on the information and returns channels without any restrictions on individuals' feasible support of ρ. As we can see, without any restrictions, the bounds are wide, ranging from the information channel share being between for -153% to 58% to the current gap. As shown in Table 1.17 in the Appendix, these upper and lower bounds are achieved when the quality of information on college earnings, R_1^2, of either Whites or Hispanics is at its highest value (60% of variance can be explained for Hispanics and 75% for Whites), and is at it's lowest value for the other group (10% of variance can be explained for Hispanics for Hispanics and 1% for Whits). As we discuss more in the next section, these extreme values are unlikely to be feasible, as they imply that individuals can explain a large share of their variance in future earnings. To make more realistic assumptions, we consider the case in which we restrict the quality of information agents have on their marginals, to be not more than a certain level. This type of restriction is similar to other restrictions which impose constraints on what variables are contained in the information set (e.g., Willis and Rosen (1979)). In our setup, the exact variables that individuals have do not matter, but what is important is how informative they are. Therefore, we take this new approach of restricting the predictive power individuals have.

Row 3 in Table 1.4 imposes a more realistic assumption on the quality of information. We restrict the share of explained variance of college attenders earnings, for both groups, to be less than 30%. Under this constraint, we can see that the information channel now contributes between -100% to -66% of the gap, which is close to our main results. Row 4 considers that individuals may explain up to 50% of the variance of their college earnings, and we can see that the bounds we get are wider, but still, information differences contribute to reducing the choice gap.

Next, we take a different approach, by relaxing the assumption that the prior mean is at the center of the support, and allow it to be anywhere on the feasible support. Row 5 in Table 1.4 shows the lower and upper bounds of the information channel and returns channel in the case where we only restrict the beliefs to lie on the feasible set. As we can see, these bounds are extremely wide, allowing for the information channel to go from reducing the gap by 16.7 percentage points (209% of the current gap) to increasing it by 10 percentage points (126% of the current gap size), where the returns channel can go from reducing the gap by 24 percentage points to increasing it by 2 percentage points.

As seen in Table 1.17, these two extreme bounds are achieved in the case where the beliefs are degenerate and are the polar opposite of the feasible set of ρ's support. Specifically, the lower bound of the decision to go to college is when Whites' $\mu_{\rho,Hispanics} = 0.89$ and Hispanics' $\mu_{\rho,Whites} = -0.9$, and the upper bound where the two are switched. This implies that the bounds are achieved in the case where the two groups have very different beliefs on how earnings behave, and they face no model uncertainty on the correlation between potential outcomes. We, therefore, consider a more realistic case, where we allow the mean priors to differ between the two groups by a small amount. Figure 1.8 shows the information channel weight and size for different mean priors. First, the black line shows the case where the two correlation values are the same. We can see that in this case, information only contributes to decreasing the gap. For most values of $\mu_{\rho,Whites}$, we observe the gap increasing by approximately 7-8 percentage points, suggesting that providing Whites with the information quality available to Hispanics would double the size of the gap, similar to our main result. The shaded area around the black line considers the case where we allow the mean beliefs values to differ between the two groups by a certain amount. The red shaded area considers the case where the differences between beliefs are allowed to differ by no more than 0.05. For all values considered, the gap suggests that reducing the information disparities between Hispanics and Whites contributes to narrowing the choice gap, thereby contributing to equality in choice. The other shaded regions show how the bounds widen as we allow the values to differ by up to 0.15.

The results in this section suggest that disparities in information quality between Whites and

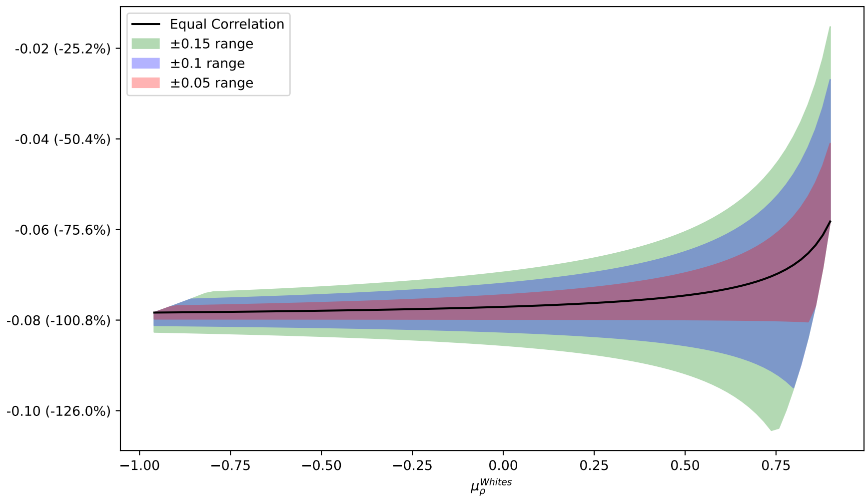

Figure 1.8: Information Channel for Different μ_ρ

Hispanics contribute to narrowing the choice gap, with a significant portion of this gap attributed to differences in returns. Should the disparities in information quality be eliminated, it would lead to an approximate 50-80% increase in the choice gap.

1.6 Assessing the Impact of Additional Information on Narrowing the Choice Gap

In the previous section we found that information differences contribute to reducing the choice gap between Hispanics and Whites. In this section we ask how can a policy maker use information in order to close the choice gap between Hispanics and Whites. We consider a policy maker with access to some information in the form of database that includes information on students characteristics, demographics, test scores, other relevant information and students outcomes labor market outcomes. This policy maker can then use this information and provide additional signal for students to better inform them on their consequences of their choices. We take the extreme case where the policy maker provide additional information only for Hispanics and ask how accurate should that information be in order to achieve parity in choice. Here, "additional information" refers to new signals that are orthogonal to an agent's existing information set; that is, we focus exclusively on previously unknown information that a policymaker could introduce. In practice, a policymaker is likely to disseminate information that correlates with what individuals already know, potentially overlapping with their private information. Therefore, in our thought exercise, we consider the case in which individuals first residualize the policymaker's signal and use only their existing information and the additional residualized information to inform their beliefs. We then examine how the information quality of this additional information affects the choice gap. Specifically, we engage in three thought exercises for this purpose. First, we consider providing information that is solely informative about earnings if the individual chooses to go to college. Second, we examine the opposite scenario where the additional information is informative only about earnings if they opt not to go to college. Finally, we consider providing information that is relevant to both types of earnings. In our thought experiments, we assume that the policymaker, akin to an econometrician, can only provide information on the marginal distributions of U_1 and U_0, as she cannot know their joint distribution. For example, the policymaker could offer students a series of tests, then provide predictions on potential earnings depending on whether they attend college or not. To measure the precision of this additional new information, we quantify it by its ability to explain the marginals of U_1 and U_0, therefore we describe these additional signals in terms of R^2 on the marginals.

To formally introduce the idea of *new information*, let s_n be the additional signal that a policymaker provides to Hispanics, after it has been partialled out from the agent's existing information. We assume that the signals are drawn from a Gaussian distribution and are correlated with α_1 and

α_0. The fact that the signal is partialled out implies that $s_n \perp\!\!\!\perp \boldsymbol{S}$, i.e. we assume that this new information is information that individuals were not able to predict given their current information set. Furthermore, as the signals and state are jointly Gaussian, the agent's beliefs are additive. Specifically, we can write the individuals' counterfactual beliefs, given their current signals and the additional information, as

$$E[U_1 - U_0|\mathcal{S}, s_n] = E[U_1 - U_0|\mathcal{S}] + \frac{\text{Cov}(U_1 - U_0, s_n)}{\text{Var}(s_n)} s_n.$$

where we used the linearity of the Gaussian distribution and the independence assumption. Next, to understand how the additional information affects choice we need to derive the variance of the counterfactual beleifs distribution. Denote by $R^2_{1,n}$ and $R^2_{0,n}$ the information quality of the new signals on U_1 and U_0, respectively. Thenwe can express the additional component in the variance of beliefs as follows:

$$\text{Var}\left(\frac{\text{Cov}(s_n, U_d)}{\text{Var}(s_n)} s_n\right) = \frac{\text{Cov}^2(s_n, \alpha_d)}{\text{Var}(s_n)} = \sigma_d^2 R^2_{d,n}$$

Without loss of generality, we can fix $\text{Var}(s_n) = 1$ and then set $\text{Cov}(s_n, U_d)^2$ to meet the required R^2_d value. Then, using the fact that $s_n \perp\!\!\!\perp \boldsymbol{S}$, we can derive the variance of the counterfactual beliefs, with the additional information:

$$\begin{aligned}\text{Var}(\text{E}[U_1 - U_0|\boldsymbol{S}, s_n]) &= \text{Var}\left(\text{E}[U_1 - U_0|\boldsymbol{S}]\right) + \text{Cov}^2(U_1, s_n) + \text{Cov}^2(U_0, s_n) - 2\text{Cov}(U_1, s_n)\text{Cov}(U_0, s_n) \\ &= \text{Var}(\text{E}[U_1 - U_0|\boldsymbol{S}]) + \sigma_1^2 R^2_{1,n} + \sigma_0^2 R^2_{0,n} - 2\sqrt{R^2_{1,n} R^2_{0,n}}\sigma_1\sigma_0.\end{aligned} \tag{1.6}$$

Given the cost function and $\mu_{\mathcal{R}}$, we can calculate the counterfactual share of students who would attend college if they were provided with this additional new information. Notice that in order to calculate the counterfactual shares we do not need to know the correlation between U_1 and U_0, as we consider how the new information is informative on the marginals, but not on the difference.

1.6.1 The Effect of Additional Information

We start by focusing on adding information exclusively to either U_1 or U_0, but not both. To achieve parity, we consider additional information previously unknown to the agent about earnings if he opts for college, which necessitates that the quality of this signal be at $R^2 = 24\%$. This implies that the additional information must independently explain almost 25% of the variance in U_1. In a similar vein, for a signal on U_0 that aims to achieve parity in choices between Whites and Hispanics, it must be capable of explaining 49% of the total variance in U_0.

Figure 1.9 explores further the counterfactual college attendance changes rate for different quality levels of additional information on U_1 and U_0, as quantified by $R^2_{1,n}$ and $R^2_{0,n}$. This figure illustrates that focusing the information predominantly on one outcome tends to enhance participation more effectively than offering a signal informative about both U_1 and U_0. This is due to the fact that information on both U_1 and U_0 reduces the variance in beliefs, as shown in equation 1.6.

Can policymakers achieve the level of accuracy as discussed above? Our analysis, detailed in Table 1.2, shows the proportion of earnings variance explained for Hispanic and White groups using our administrative data. The table presents out-of-sample R^2 values from an Extreme Gradient Boosting model, which predicts earnings 12-14 years post high school graduation. This model incorporates students' characteristics, exit exam scores, course selections in high school, and pass-fail for each course, for both Hispanics and Whites. Such an analysis simulates the data a policy maker might want to utilize in advising students about college decisions. The table shows that approximately only 10% of the variance in earnings for both college attendees and non-attendees in our sample can be explained using this information. This is much lower than the needed level of information quality to achieve equality of choice. Introducing fixed effects for schools into the model does not markedly improve prediction accuracy. We preform a similar exercise using the National Longitudinal Survey of Youth 1997 (NLSY97), as shown in Table 1.16 in the Appendix. Here, due to to a smaller sample size, we employed linear regression to estimate earnings for individuals aged 34 or 35, both college attendees and non-attendees. The NLSY97 dataset provides extensive individual data, covering aspects like gender, cohort, urbanicity, abilities (measured via ASVAB tests), parental education and income, and high school performance. Most importantly, the direct measures of ability and parental income, are typically not available in administrative datasets, thus set a potential upper limit on the prediction quality a policy maker can make. Our analysis using adjusted R^2[10] reveals that up to 17% of earnings variance for non-college Hispanic and and less than 10% for other groups, can be accounted for. It's important to note that in our counterfactual exercise above, we consider providing *new* information to students. A large share of the information schools can provide is already known to students and, therefore, is even less likely to generate significant changes in behavior.

Other research has noted the limitations of current models, measurements, and approaches in explaining variations in outcome variables of interest in social science (Salganik et al. (2020),Garip (2020)). Specifically, similar to our study, other papers examined how different pre-college measurements of ability, such as IQ, achievement tests, high school grades, or personality tests, explain the variance in earnings and other metrics (Murnane et al. (2000), Watts (2020), Borghans et al. (2016)). They found that these measurements explain up 20%.[11] These results collectively suggest that our

10. We use in-sample adjusted R^2 due to the smaller sample size. This approach, while not ideal, is frequently employed in literature discussing the prediction of earnings based on high school performance (Murnane et al. (2000),Watts (2020),Borghans et al. (2016)).

11. It's important to note that these measures usually use in-sample R^2, or adjusted R^2. which are usually higher than the out-of-sample ones (Hastie et al. (2009)). The out-sample R^2 is the one we care about as it do not suffer the over fitting

standard data, which are likely to use in any recommendation systems for college, is not informative enough in order to explain future earnings. Therefore, achieving equality through informational interventions might be challenging.

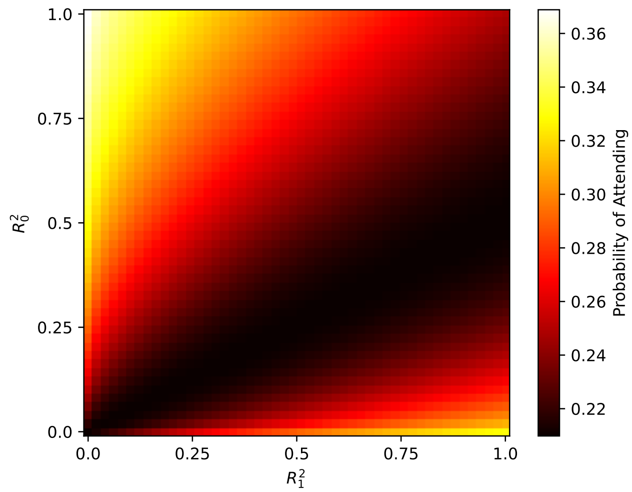

Figure 1.9: The effect of additional information on Earnings

Note: Figure 1.9 shows the counterfactual share of Hispanics who would attend college after providing them with an additional signal of information quality on U_1 of R_1^2 and information quality on U_0 of R_0^2. For both figures, the quality of information is measured based on the ability to explain quarterly earnings 12-15 years after high school graduation.

1.7 Conclusions

Individuals from diverse backgrounds have unique upbringings that significantly shape their later life and subsequent choices. Such experiences are pivotal in defining the constraints and opportunities they encounter, along with the outcomes and their information on these outcomes. This project explores how differences in returns and information affect different groups' college-going decisions.

In this context, differences in the outcomes and information can be driven by disparities that take place prior to the time the decision is made (Neal and Johnson (1996)). For example, information differences can arise when affluent groups have access to better information on college outcomes compared to less fortunate ones. Similarly, growing up in wealthier backgrounds can also lower college costs for children. More generally, differences in information and potential returns are likely to rise in dynamic models, where past decisions affect the current decision environment. For instance, Cunha and Heckman (2007a) and Cunha et al. (2021) illustrate three ways through which early life disparities can shape future opportunities: through affecting the choice set, the dynamic incentives, and the information individuals possess about the outcomes of investments. Better understanding how these past disparities affect future disparities is crucial to understand what drives disparities.

In this project, we focused on the gaps in the quality of information. Differences in quality are not solely a byproduct of past decisions and disparities but are also driven by the future. This is particularly evident when considering the challenges associated with predicting labor market outcomes. Earnings for different groups can vary widely due to factors like industry sector trends, geographic economic conditions, and social biases. These disparities affect not only the returns distribution but also the ability to predict future returns. For example, minorities may suffer from discrimination that results in them earning lower wages, but this outcome is easier to predict than the case where some earn low wages and some high wages. If there is less uncertainty in future earnings, it may be easier to predict the future. Therefore, focusing on information quality allows measures through which both the past and the future affect current disparities.

This paper introduces a new approach to analyze how differences in information and potential returns across groups impact choice disparities. Therefore, we take a more systematic approach (Bohren et al. (2022),Small and Pager (2020)) not focusing on a specific channel or variable that affects choice but focusing on the cumulative effect of past and future disparities on the current choice gap. In our empirical exercise, we find that the information differences between Hispanics and Whites help to mitigate the choice gap, implying that future and past disparities contribute to the gap mainly through incentives and not through information on these incentives. In the second part of the paper, we find that achieving parity in choice through policy interventions that provide additional information to Hispanics may be extremely difficult, as the amount and quality of additional information Hispanics need is extremely high. This suggests that while information-based initiatives may have limited effectiveness, strategies directly targeting outcomes may be more effective in the long term to achieve parity in choices.

Finally, the approach proposed in this paper could be applied to other scenarios where it may be interesting to quantify drivers of choice gaps, such as cases of discrimination, healthcare, and other decisions related to investing in human capital and skill development. The central idea we present here is that in order to comprehend the drivers of behavioral differences and choices, as well as why these disparities persist, we must put a spotlight on the information environment in which people make decisions. Understanding the informational environment in which people operate is essential for understanding why and how differences across groups are formed and persist.

1.8 Additional Figures

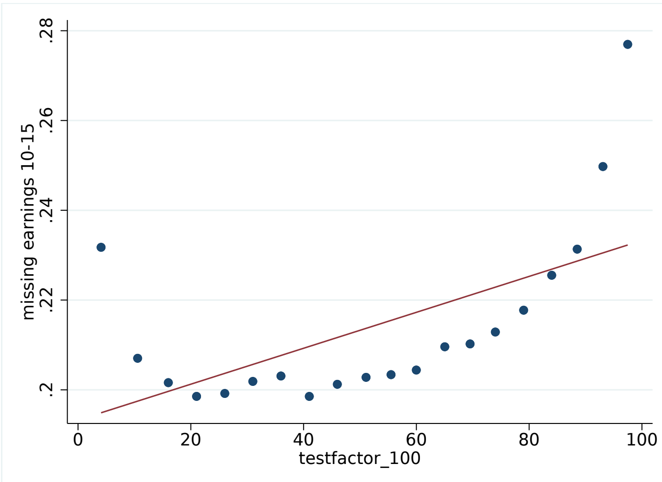

(a) Relation between test scores and missing Earnings for earnings 10-15 years after high school graduation

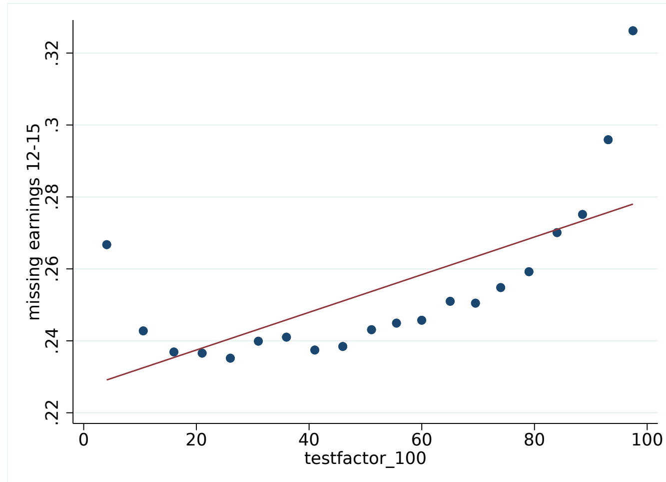

(b) Relation between test scores and missing Earnings for earnings 12-15 years after high school graduation

Figure 1.10: Relation Between test scores and Missing Earnings

Notes: The above figures plot the share of missing earnings by test score factor, as described in Section 1.3. The first figure presents the missing earnings for the period of 10-15 years after high school graduation. Figure (b) illustrates the share of missing earnings for the period of 12-15 years after high school graduation. The red line indicates the expected trend line.

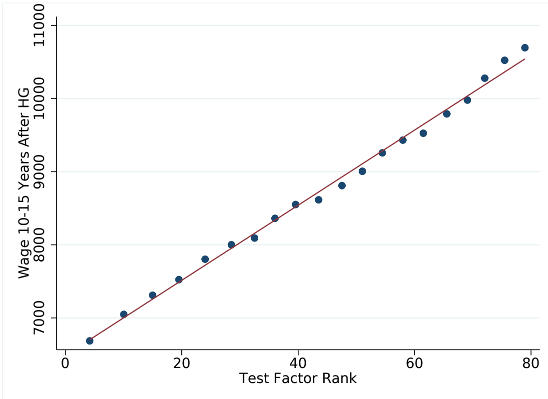

(a) Relation between test scores and earnings

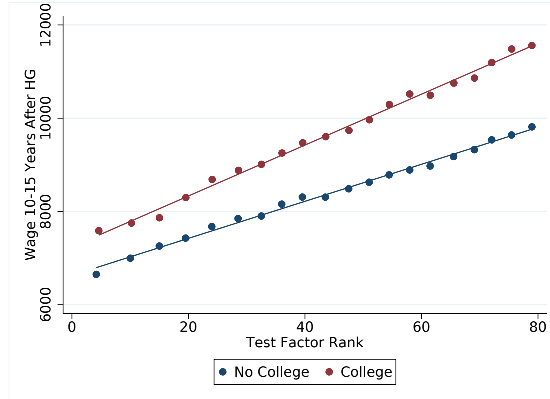

(b) Relation between test scores and earnings, by college attendance

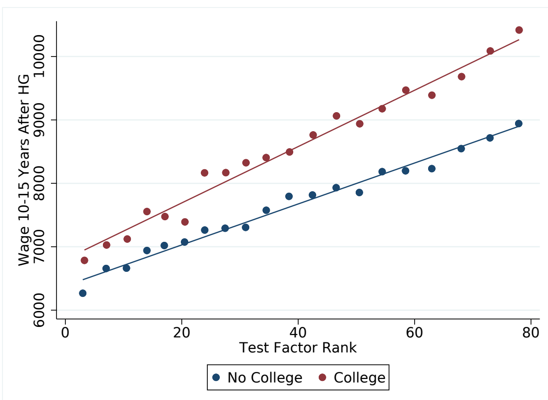

(c) Relation between test scores and earnings, by college attendance - Hispanics

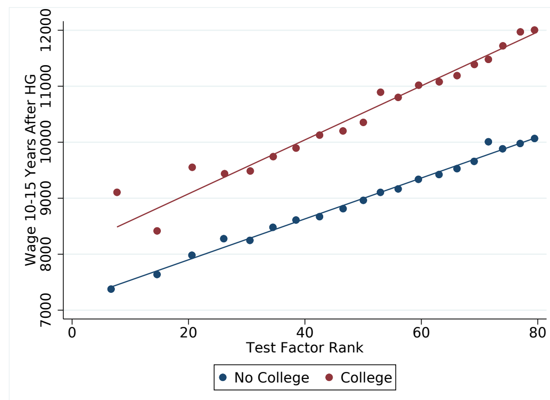

(d) Relation between test scores and earnings, by college attendance - Whites

Figure 1.11: Relation Between Test Scores and Earnings

Notes: This figure illustrates the relationship between test score percentile, as calculated in Section 1.3, and the expected average earnings 10-15 years after high school graduation. Figure (a) depicts the correlation between test scores and earnings for all individuals. Figure (b) presents this relationship, separated for individuals who attended college (red line) and those who did not (blue line). Figure (c) displays the same data but specifically for Hispanic individuals, while figure (d) focuses exclusively on White individuals.

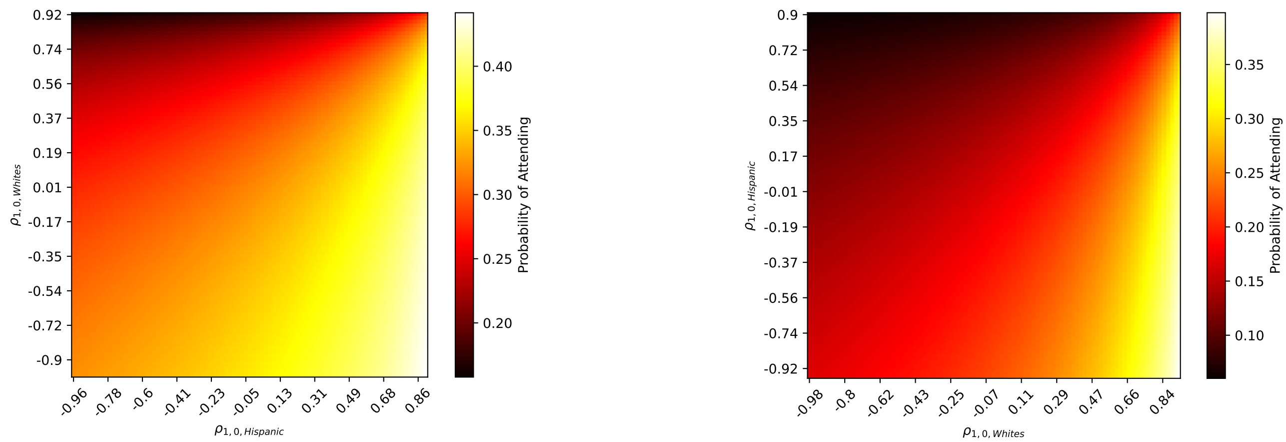

(a) Counterfactuals share of Whites with Hispanic Information

(b) Counterfactuals share of Hispanics with Whites Information

Figure 1.12: Counterfactuals share

Notes: The figures illustrate the counterfactual share of White and Hispanic college attendance for various potential earnings correlation values. Figure (a) depicts the share of White individuals under the scenario where they are provided with the same quality of information as Hispanics, with information quality measured across different correlation values. Figure (b) presents the counterfactual shares of Hispanic college attendance, assuming they received the information quality of Whites, as gauged by varying correlation values.

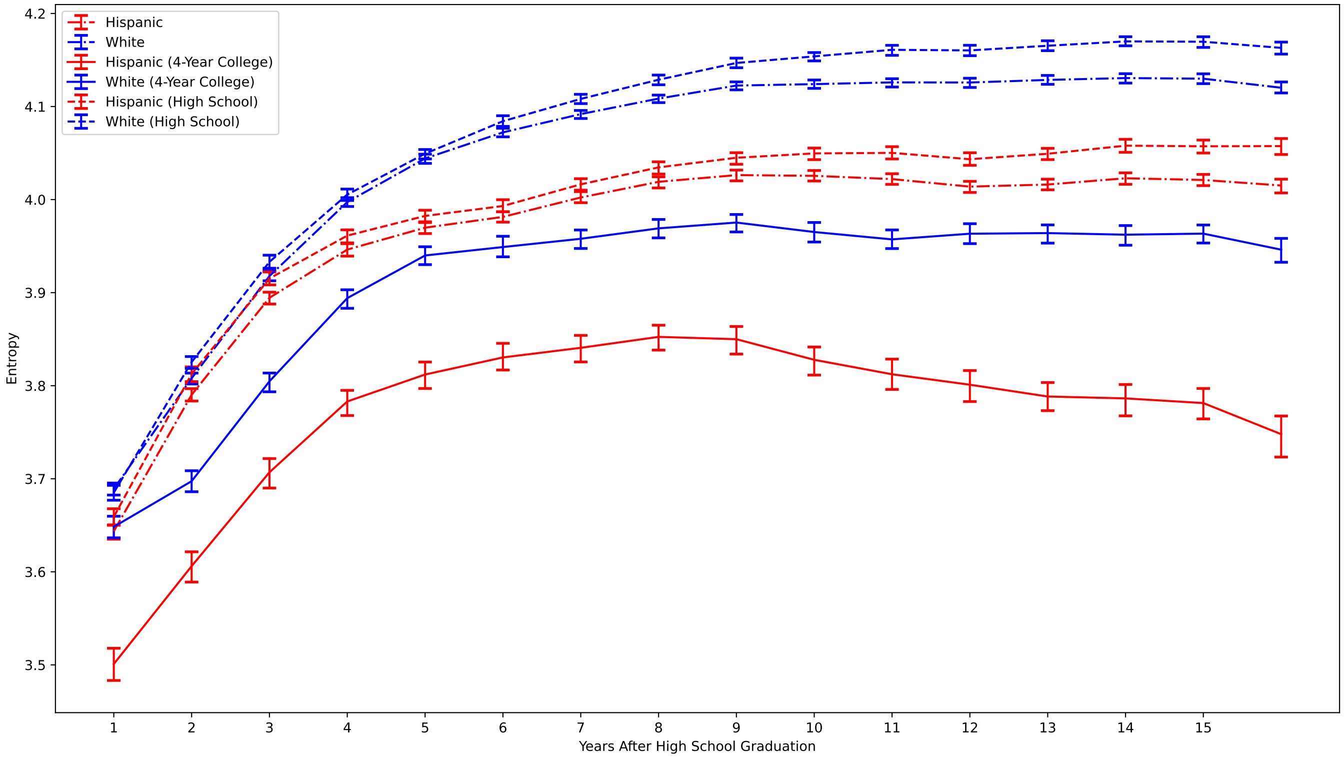

Figure 1.13: Shannon's Entropy for NAICS Industries

Notes:The figure displays the entropy of 2-digit NAICS code industries in which Hispanics and Whites, who attended 4 year college and who did not, are employed, plotted against the number of years post-high school graduation on the x-axis. Confidence intervals are at 95%.

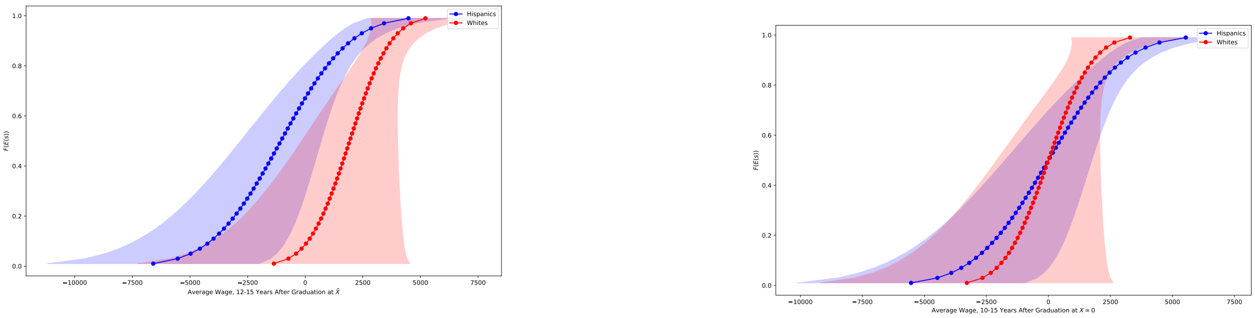

(a) Counterfactuals share of whites with Hispanic Information

(b) Beliefs CDF - Demean

Figure 1.14: Beliefs Cumulative Distribution for Whites and Hispanics

Note: The figure displays the Cumulative Distribution Function of beliefs on α_1 *and* α_0 *for both Hispanics and Whites. The shaded area represents the 95% Confidence Interval. Figure (a) illustrates these beliefs for the case where covariates are set to their mean. Figure (b) depicts the same graphs with all covariates, including the constant, set to zero.*

1.9 Additional Tables

	All	Hispanic	Whites
College Attendance	0.23 (0.42)	0.18 (0.38)	0.27 (0.44)
Test Factor Percentile	43.18 (22.02)	36.37 (22.02)	48.11 (20.67)
Math Score	45.62 (23.88)	40.29 (24.12)	49.49 (22.94)
Reading Score	47.5 (25.41)	40.11 (25.41)	52.87 (24.03)
No Disadvantage	0.7 (0.46)	0.41 (0.49)	0.91 (0.29)
Elig. Free Meals	0.22 (0.41)	0.44 (0.5)	0.06 (0.24)
Elig. Reduced Price Meals	0.06 (0.23)	0.09 (0.29)	0.03 (0.16)
Other Disadvantage	0.03 (0.16)	0.06 (0.23)	0.0 (0.05)
Distiguish	0.06 (0.24)	0.07 (0.25)	0.05 (0.23)
Minimal	0.22 (0.41)	0.19 (0.39)	0.24 (0.43)
Required	0.72 (0.45)	0.74 (0.44)	0.7 (0.46)
CT Median Income	44027.0 (21371.0)	36265.0 (15939.0)	49663.0 (22986.0)
CT Families Below Poverty Line	14.5 (10.82)	20.08 (12.19)	10.44 (7.42)
CT Share of Employed	63.21 (9.97)	59.92 (10.01)	65.6 (9.23)
Title I schools	0.34 (0.47)	0.58 (0.49)	0.17 (0.38)
No Participation in Tech Program	0.24 (0.43)	0.22 (0.41)	0.26 (0.44)
Enroll in Career Tech Elective (6-12)	0.23 (0.42)	0.2 (0.4)	0.24 (0.43)
Participate in Tech Prep Prog (9-12)	0.32 (0.47)	0.33 (0.47)	0.32 (0.47)
Participate in Tech Prep Prog	0.21 (0.41)	0.25 (0.43)	0.18 (0.38)
Share in Oil Industry	52.73 (28.53)	49.21 (29.14)	55.29 (27.79)
City	0.37 (0.48)	0.52 (0.5)	0.25 (0.44)
Suburb	0.32 (0.47)	0.24 (0.43)	0.38 (0.49)
Town	0.11 (0.31)	0.11 (0.31)	0.1 (0.31)
Rural	0.2 (0.4)	0.13 (0.34)	0.26 (0.44)
Distance to 4-Year College	19.82 (18.8)	18.19 (20.5)	21.04 (17.25)

Table 1.5: Summary Statistics

Note: The Columns include 12th-grade analysis cohorts from 2003-2005. NCES geographic categories are condensed into four types (city, suburb, town, rural). Distance from College is measured using the geodesic distance from the student high school to near by college. CT stands for the School Census Tract. Distinguish, minimal and required are the share of studnets with the Distinguished Achievement Program, Recommended High School Program, or the Minimum High School Program, respectivly. College Attendacne capture the share of high school students who attended college in the first year after high school graduation year

	All	Hispanic	Whites	Difference (Whites - Hispanic)
Wage 8-10	7117.0 (4533.0)	6393.0 (3974.0)	7627.0 (4823.0)	1234.0 (6249.3)
Wage 10-12	8215.0 (5194.0)	7348.0 (4509.0)	8852.0 (5558.0)	1504.0 (7157.0)
Wage 12-14	9079.0 (5808.0)	8046.0 (4952.0)	9823.0 (6249.0)	1777.0 (7973.2)
Wage 14-16	9838.0 (6280.0)	8721.0 (5383.0)	10658.0 (6748.0)	1937.0 (8632.0)
Wage 12-15	9214.0 (5807.0)	8209.0 (4993.0)	9959.0 (6239.0)	1750.0 (7990.9)

Table 1.6: Wages Summary Statistics

Note: The table presents the mean earnings for Hispanics and Whites across various periods, spanning 8-16 years after high school graduation. For the period of 14-16 years post-graduation, data is exclusively from the 2003-2004 cohort. For all other time frames, data includes all cohorts from 2003-2004. Standard deviations are provided in parentheses.

	Hispanics		Whites	
	No College	College	No College	College
Wage 1-2	2807.0	1904.0	2717.0	1693.0
	(1785.0)	(1351.0)	(1910.0)	(1340.0)
Wage 3-4	3903.0	2903.0	3935.0	2862.0
	(2465.0)	(2070.0)	(2808.0)	(2248.0)
Wage 5-7	4880.0	5027.0	5388.0	5983.0
	(2984.0)	(3128.0)	(3550.0)	(3693.0)
Wage 8-10	6234.0	7238.0	7237.0	8775.0
	(3973.0)	(4209.0)	(4757.0)	(4961.0)
Wage 10-12	7066.0	8468.0	8299.0	10285.0
	(4403.0)	(4762.0)	(5364.0)	(5808.0)
Wage 12-14	7750.0	9424.0	9201.0	11452.0
	(4804.0)	(5258.0)	(5918.0)	(6512.0)
Wage 14-16	8360.0	10180.0	9973.0	12447.0
	(5236.0)	(5736.0)	(6441.0)	(7211.0)
Wage 12-15	7862.0	9596.0	9327.0	11618.0
	(4849.0)	(5325.0)	(5978.0)	(6616.0)

Table 1.7: Wages Summary by College Statistics

Note: The table presents the mean and standard deviation earnings for Hispanics and Whites across various periods after high school graduation For the period of 14-16 years post-graduation, data is exclusively from the 2003-2004 cohort. For all other time frames, data includes all cohorts from 2003-2004. Standard deviations are provided in parentheses.

	All	Hispanic	Whites
No Controls	-0.0156	-0.0232	-0.0436
	(0.0054)	(0.0053)	(0.0038)
Ind. Controls	-0.0277	-0.0151	-0.0392
	(0.0044)	(0.0066)	(0.0045)
+ School Char.	-0.0061	0.0074	-0.0177
	(0.004)	(0.0057)	(0.004)
+ Neighborhood Char.	-0.0014	0.0009	-0.0036
	(0.0018)	(0.0022)	(0.0021)

Table 1.8: Instrument Diagnostics

Note: The table displays coefficients on distance to a 4-year college, derived from a regression of test score factors, as defined in section 1.3, on distance to college. Each row introduces additional controls for individual student characteristics, school characteristics, and neighborhood characteristics. Standard errors, provided in parentheses, are clustered at the school-cohort level.

	All	Hispanic	Whites
No Controls	-0.0008	-0.0007	-0.0013
	(0.0001)	(0.0002)	(0.0001)
	317278	136581	180697
Ind. Controls	-0.0008	-0.0006	-0.0011
	(0.0001)	(0.0001)	(0.0001)
	317278	136581	180697
+ School Char.	-0.0014	-0.001	-0.0019
	(0.0002)	(0.0002)	(0.0002)
	317278	136581	180697
+ Neighborhood Char.	-0.0016	-0.0023	-0.0012
	(0.0002)	(0.0003)	(0.0002)
	317278	136581	180697

Table 1.9: First Stage

Note: The table presents the first-stage regression results, analyzing the effect of distance to a 4-year college on college attendance in the first year post-graduation. Each row adds additional controls for individual student characteristics, school characteristics, and neighborhood characteristics, as defined in section 1.3. Standard errors are given in parentheses and are clustered at the school-cohort level.

	All				Hispanic				Whites			
Wage Avg	8-10	10-12	12-14	14-16	8-10	10-12	12-14	14-16	8-10	10-12	12-14	14-16
No Controls	7.2078	3.1346	-0.6775	-3.4851	6.7659	4.1617	1.9655	1.4303	3.362	-3.329	-9.8946	-15.7508
	(1.5094)	(1.4327)	(1.4515)	(1.8191)	(1.5053)	(1.2071)	(1.1988)	(1.5494)	(1.2177)	(1.4513)	(1.6605)	(2.2097)
Obs.	245206	239307	233091	149498	103198	101284	99428	63271	142008	138023	133663	86227
Ind. Controls	4.9314	0.5937	-3.5101	-6.5359	5.8931	3.3429	1.2646	0.7636	4.1198	-2.1051	-8.4963	-14.2128
	(1.0258)	(0.946)	(1.0365)	(1.4211)	(1.4953)	(1.2306)	(1.1954)	(1.5131)	(1.1556)	(1.3397)	(1.5095)	(1.9959)
Obs.	245206	239307	233091	149498	103198	101284	99428	63271	142008	138023	133663	86227
+ School Char.	-1.7881	-2.9513	-4.1313	-5.7861	-1.7207	-1.0352	-0.8267	-0.6872	-1.758	-4.7777	-7.6157	-11.5689
	(1.2276)	(1.2815)	(1.3991)	(1.8907)	(1.384)	(1.4014)	(1.5289)	(2.0046)	(1.4649)	(1.7076)	(1.8884)	(2.5429)
Obs.	245206	239307	233091	149498	103198	101284	99428	63271	142008	138023	133663	86227
+ Neighborhood Char.	-0.504	-1.7866	-3.2171	-5.9322	-1.9984	-0.8557	-0.7281	-1.1675	1.6792	-0.7804	-3.5928	-8.7372
	(1.6783)	(1.9833)	(2.2052)	(2.8903)	(2.4266)	(2.8765)	(3.114)	(4.094)	(2.1061)	(2.4762)	(2.7454)	(3.6005)
Obs.	245206	239307	233091	149498	103198	101284	99428	63271	142008	138023	133663	86227

Table 1.10: Reduced Form

Note: The table presents the reduced-form results of regressing the distance to a 4-year college on earnings for the periods 8-10, 10-12, and 14-16 years after high school graduation. Each row incorporates additional controls for individual student characteristics, school characteristics, and neighborhood characteristics, as defined in Section 1.3. For all periods, the data includes the three cohorts from 2003-2005. Specifically for the 14-16 year period, only the 2003-2004 cohorts are used. Standard errors, provided in parentheses, are clustered at the school-cohort level.

	Hispanics Coefficient		
0	Baseline	-0.0891	(0.0092)
1	+ Neighborhood Char.	-0.1317	(0.0046)
2	+ Individual Chars.	-0.0867	(0.0039)
3	+ School Char.	-0.0776	(0.0038)
4	+ Test Score	-0.0428	(0.0035)

Table 1.11: College Attendance Gap

Note: The table displays the coefficient for Hispanics from a regression analysis, where the dependent variable is an indicator of first-time college attendance and the independent variable is the indicator of being Hispanic. Each row adds additional controls. The first row represents the raw difference with a cohort fixed effect. Subsequent rows include additional controls for individual student characteristics, school characteristics, and neighborhood characteristics, as defined in Section 1.3. Standard errors are shown in parentheses and are clustered at the school-cohort level.

		Grades X College	R2	N
Wage 12-14	All	15.88	0.17	236092
		(1.35)		
Wage 12-14	Hispanics	12.59	0.15	100140
		(1.89)		
Wage 12-14	Whites	14.98	0.15	135952
		(1.92)		
Wage 14-16	All	16.16	0.18	151336
		(1.8)		
Wage 14-16	Hispanics	11.26	0.16	63734
		(2.58)		
Wage 14-16	Whites	15.56	0.16	87602
		(2.61)		
Wage 12-15	All	15.74	0.17	240692
		(1.36)		
Wage 12-15	Hispanics	12.74	0.15	101854
		(1.91)		
Wage 12-15	Whites	14.86	0.16	138838
		(1.89)		

Table 1.12: Relation Between Earnings and Grades

Note: The table displays the coefficient on the interaction term for Exit Exam Grades and College Attendance in the first year after high school graduation. Standard errors, presented in parentheses, are clustered at the school-cohort level.

	All	Hispanics	Whites
Ind Chr.	0.7	0.68	0.67
+ School Char.	0.72	0.71	0.71
+ Neighberhood Char.	0.75	0.75	0.74
+ Test Scores	0.78	0.8	0.77
N	321411	137551	183860

Figure 1.15: Area Under the Curve Analysis of Predicting College Attendance Decisions

Note: This table presents the Area Under The Curve (AUC) from a Probit model, predicting college attendance in the first year after high school among graduates. Each row progressively includes additional controls. The first row incorporates individual characteristics, the second includes school characteristics, the fourth integrates neighborhood characteristics, and the final row additionally accounts for the test score factor. For a detailed description of these controls, refer to section 1.3.

	school matters?		personal matters?		career options?		college options?		high school rank?		Top 10% rule?	
	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites
Yes	0.65	0.63	0.26	0.23	0.56	0.45	0.61	0.58	0.48	0.54	0.26	0.33
	(0.006)	(0.007)	(0.006)	(0.006)	(0.006)	(0.007)	(0.006)	(0.007)	(0.007)	(0.007)	(0.006)	(0.007)
No	0.21	0.19	0.42	0.36	0.31	0.33	0.29	0.26	0.40	0.33	0.55	0.45
	(0.005)	(0.006)	(0.006)	(0.007)	(0.006)	(0.007)	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)	(0.007)
Have not needed	0.13	0.17	0.31	0.41	0.12	0.22	0.09	0.15	0.10	0.12	0.17	0.21
	(0.004)	(0.006)	(0.006)	(0.007)	(0.004)	(0.006)	(0.004)	(0.005)	(0.004)	(0.005)	(0.005)	(0.006)
No response	0.01	0.01	0.00	0.00	0.01	0.01	0.01	0.01	0.01	0.01	0.02	0.01
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
Weighted Obs	139644	212774	139644	212774	139644	212774	139644	212774	139644	212774	139644	212774
Unweighted Obs	11114	12621	11114	12621	11114	12621	11114	12621	11114	12621	11114	12621

Table 1.13: Sources of Information

Note: This table displays responses from Senior and Sophomore cohorts participating in the Texas Higher Education Opportunity Project regarding the dissemination of information by school counselors on the subjects indicated in the header. It quantifies the proportions of Hispanic and White students who answered "Yes," "No," "Not needed," or did not respond. Standard errors are in parentheses.

	course selection		personal problems		school discipline		jobs		educational plans		choosing a college		college applications		letters of rec.		college essays		financial aid		job interviews	
	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites
Three or more times	0.14	0.10	0.02	0.02	0.01	0.01	0.01	0.01	0.08	0.07	0.10	0.07	0.10	0.09	0.06	0.05	0.03	0.03	0.07	0.04	0.01	0.00
	(0.005)	(0.004)	(0.002)	(0.002)	(0.001)	(0.001)	(0.002)	(0.002)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.003)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)	(0.001)	(0.001)
Twice	0.29	0.28	0.04	0.03	0.04	0.02	0.04	0.03	0.13	0.12	0.07	0.07	0.07	0.07	0.04	0.05	0.03	0.03	0.05	0.04	0.01	0.01
	(0.006)	(0.007)	(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)	(0.004)	(0.005)	(0.003)	(0.004)	(0.003)	(0.004)	(0.002)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)	(0.001)	(0.001)
Once	0.28	0.32	0.09	0.08	0.08	0.05	0.13	0.12	0.23	0.26	0.10	0.11	0.11	0.12	0.07	0.09	0.06	0.06	0.11	0.10	0.02	0.01
	(0.006)	(0.007)	(0.004)	(0.004)	(0.004)	(0.003)	(0.004)	(0.005)	(0.006)	(0.006)	(0.004)	(0.005)	(0.004)	(0.005)	(0.003)	(0.004)	(0.003)	(0.004)	(0.004)	(0.004)	(0.002)	(0.002)
Never	0.19	0.21	0.82	0.85	0.84	0.89	0.80	0.81	0.50	0.51	0.15	0.19	0.14	0.17	0.25	0.26	0.30	0.33	0.18	0.25	0.38	0.42
	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.006)	(0.007)	(0.007)	(0.005)	(0.006)	(0.005)	(0.006)	(0.006)	(0.006)	(0.006)	(0.007)	(0.005)	(0.006)	(0.006)	(0.007)
No response	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.00
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Weighted Obs	139644	212774	139644	212774	139644	212774	139644	212774	139644	212774	139644	212774	139644	212774	139644	212774	139644	212774	139644	212774	139644	212774
Unweighted Obs	11114	12621	11114	12621	11114	12621	11114	12621	11114	12621	11114	12621	11114	12621	11114	12621	11114	12621	11114	12621	11114	12621

Table 1.14: Number of Interactions with School Councilor

Note: This table presents responses from Senior and Sophomore cohorts involved in the Texas Higher Education Opportunity Project, detailing the frequency of their interactions with the school counselor in the past year regarding the topics listed in the header. It quantifies the proportions of Hispanic and White students who indicated their interactions as "Three or more times," "Twice," "Once," "Never," or did not respond. Standard errors are provided in parentheses.

	Education		Important Issues		Job		Relationships		Finance	
	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites	Hispanic	Whites
Often	0.42	43.00	0.39	47.00	0.39	37.00	0.36	28.00	0.27	43.00
	(0.50)	(0.36)	(0.49)	(0.37)	(0.49)	(0.31)	(0.48)	(0.22)	(0.45)	(0.34)
Sometimes	0.44	58.00	0.39	61.00	0.43	67.00	0.51	67.00	0.56	54.00
	(0.50)	(0.49)	(0.49)	(0.48)	(0.50)	(0.56)	(0.51)	(0.53)	(0.50)	(0.42)
Never	0.13	18.00	0.22	20.00	0.17	16.00	0.13	32.00	0.18	31.00
	(0.34)	(0.15)	(0.42)	(0.16)	(0.38)	(0.13)	(0.34)	(0.25)	(0.39)	(0.24)
Num Obs	58.00	144	54.00	155	56.00	141	55.00	123	49.00	140

Table 1.15: Number of Interactions with Parents

Note: This table presents data from the National Longitudinal Survey of Youth 1997 (NLSY97), focusing on participants' responses to questions regarding the frequency with which they discuss the topics listed in the header with their mother or father. It illustrates the proportion of times participants reported discussing these topics "often" with at least one parent, "sometimes" with at least one parent, or "never" with both parents. Standard errors are provided in parentheses.

	Baseline			Ability			Ability + Parental Income			Ability + Parental Income + Parental Educ			Ability + Parental Income + Parental Educ+ Grades		
	R^2	$R^2 - Adj.$	N	R^2	$R^2 - Adj.$	N	R^2	$R^2 - Adj.$	N	R^2	$R^2 - Adj.$	N	R^2	$R^2 - Adj.$	N
All	0.139	0.137	3568.0	0.153	0.150	2965.0	0.163	0.159	2092.0	0.158	0.151	1490.0	0.135	0.117	910.0
Whites	0.129	0.126	2554.0	0.137	0.134	2192.0	0.150	0.145	1578.0	0.147	0.138	1185.0	0.123	0.102	749.0
Hispanics	0.131	0.125	1014.0	0.183	0.174	773.0	0.204	0.189	514.0	0.234	0.202	305.0	0.289	0.205	161.0
Whites- No College	0.085	0.081	1404.0	0.106	0.100	1171.0	0.122	0.112	848.0	0.132	0.115	587.0	0.148	0.097	284.0
Whites - College	0.046	0.041	1150.0	0.057	0.050	1021.0	0.080	0.068	730.0	0.093	0.076	598.0	0.105	0.073	465.0
Hispanics- No College	0.083	0.076	757.0	0.116	0.105	569.0	0.145	0.124	382.0	0.180	0.135	215.0	0.300	0.172	104.0
Hispanics - College	0.048	0.025	257.0	0.113	0.081	204.0	0.104	0.038	132.0	0.177	0.061	90.0	0.250	-0.050	57.0

Table 1.16: NLSY97 - R^2

Note: This table uses data from the National Longitudinal Survey of Youth 1997 (NLSY97), focusing on students who were 16 and 17 years old in 1997, to show the prediction quality of their income in 2015 using pre-decision variables. It presents the R^2 *and adjusted* R^2 *values, from regression for All, Hispanics and whites and by college attendance. The "Baseline" column accounts for social group, gender, birth year, and college attendance. Subsequent columns incrementally introduce additional variables: the "Ability" columns include ASVAB test results; the third column incorporates household income data; the fourth column integrates information on parental education levels; and the final column incorporates high school grade information.*

	Information Channel	Returns Channel
1) $\mu_\rho = 0$	-0.077 (-97.116%)	0.156 (197.0%)
2) All Possible R_1^2		
LB, CF= 0.41	-0.121 (-152.532%)	0.2 (253.0%)
Whites Support, $R_1^2 = 0.01$	$[-0.9, 0.89]$	
Hispanics Support, $R_1^2 = 0.6$	$[0.62, 0.73]$	
UB, CF= 0.243	0.046 (57.807%)	0.033 (42.193%)
Whites Support, $R_1^2 = 0.75$	$[0.8, 0.87]$	
Hispanics Support, $R_1^2 = 0.1$	$[-0.9, 0.89]$	
3) $R_1^2 \leq 0.3$		
LB, CF= 0.369	-0.079 (-100.15%)	0.159 (200.0%)
Whites Support, $R_1^2 = 0.01$	$[-0.9, 0.89]$	
Hispanics Support, $R_1^2 = 0.24$	$[-0.56, 0.89]$	
UB, CF= 0.342	-0.053 (-66.8%)	0.132 (167.0%)
Whites Support, $R_1^2 = 0.25$	$[-0.42, 0.89]$	
Hispanics Support, $R_1^2 = 0.1$	$[-0.9, 0.89]$	
4) $R_1^2 \leq 0.5$		
LB, CF= 0.378	-0.089 (-112.35%)	0.169 (212.0%)
Whites Support, $R_1^2 = 0.01$	$[-0.9, 0.89]$	
Hispanics Support, $R_1^2 = 0.34$	$[-0.27, 0.89]$	
UB, CF= 0.325	-0.036 (-45.38%)	0.115 (145.0%)
Whites Support, $R_1^2 = 0.41$	$[-0.06, 0.89]$	
Hispanics Support, $R_1^2 = 0.1$	$[-0.9, 0.89]$	
5) Unrestricted Mean Beliefs		
LB, CF= 0.456	-0.167 (-209.918%)	0.246 (310.0%)
Whites μ_ρ	0.89	
Hispanics μ_ρ	-0.9	
UB, CF= 0.189	0.1 (126.0%)	-0.021 (-25.952%)
Whites μ_ρ	-0.9	
Hispanics μ_ρ	0.89	

Table 1.17: Main Decomposition - Extended

Note: This table provides additional information on the support and R_1^2 that achieve the bounds in Table 1.4. Row 1 shows our main results, the decomposition of the choice gap into the information channel and returns channel. Rows 2-4 show the upper bound (UB) and lower bound (LB) of the information channel under the different assumptions on the quality of information individuals have on their college earnings (R_1^2), as discussed in the main text. In each row, the table shows the support of ρ's that are induced by the R_1^2 for whites and Hispanics for both the UB and LB. The bounds on the Counterfactual (CF) share of whites who would go to college if they had the information quality of Hispanics are shown for each case. Row 5 shows the lower and upper bounds of the information channel and returns channel in the case where we only restrict the mean prior beliefs to lie on the feasible set ρ, without restriction on R_1^2. For both the LB and UB, the table shows the mean μ_ρ for each group that attained these bounds.

1.10 Identification

1.10.1 Proof of Proposition 1

In this section we prove proposition 1 in the main text text. We note that that the proof here is for the empircal specifciation, introduced in section 1.2.5, but the proof can trivially be extended to the non-parametric case outlined in 1.10.2.

Proposition 1 *Fix R_1^2. A ρ is feasible from the high school graduate perspective if and only if it is feasible from the Econometrician's perspective.*

Before proving the main proposition, we first show the following lemma. Let R_1^2 be fixed. Consider the covariance matrix $C_E(\rho)$ associated with the random variables U_1, U_0, $\mathbb{E}[U_1|\boldsymbol{S}]$, and $\mathbb{E}[U_0|\boldsymbol{S}]$. The matrix $C_E(\rho)$ is defined as:

$$C_E(\rho) = \begin{bmatrix} \sigma^2_{U_1} & \rho\sigma_{U_1}\sigma_{U_0} & \sigma_{U_1,\mathbb{E}[U_1|\boldsymbol{S}]} & \sigma_{U_1,\mathbb{E}[U_0|\boldsymbol{S}]} \\ \rho\sigma_{U_0}U_1 & \sigma^2_{U_0} & \sigma_{U_0,\mathbb{E}[U_1|\boldsymbol{S}]} & \sigma_{U_0,\mathbb{E}[U_0|\boldsymbol{S}]} \\ \sigma_{\mathbb{E}[U_1|\boldsymbol{S}],U_1} & \sigma_{\mathbb{E}[U_1|\boldsymbol{S}],U_0} & \sigma^2_{\mathbb{E}[U_1|\boldsymbol{S}]} & \sigma_{\mathbb{E}[U_1|\boldsymbol{S}],\mathbb{E}[U_0|\boldsymbol{S}]} \\ \sigma_{\mathbb{E}[U_0|\boldsymbol{S}],U_1} & \sigma_{\mathbb{E}[U_0|\boldsymbol{S}],U_0} & \sigma_{\mathbb{E}[U_1|\boldsymbol{S}],\mathbb{E}[U_0|\boldsymbol{S}]} & \sigma^2_{\mathbb{E}[U_0|\boldsymbol{S}]} \end{bmatrix}, \tag{1.7}$$

where σ^2_X denotes the variance of X, $\sigma_{X,Y}$ denotes the covariance between X and Y, and ρ is the correlation coefficient between U_1 and U_0. All elements of $C_E(\rho)$ are identified except for ρ.

Proof. The identification of σ_1 and σ_0, $\text{Var}(E[U_1 - U_0|\boldsymbol{S}])$, $\text{Cov}(U_d, E[U_1 - U_0|\boldsymbol{S}])$ are shown in the main text. We proceed in showing identification of the other components. Identification of $\sigma^2_{E[U_1|\boldsymbol{S}]}$

stems from the equality $\sigma^2_{E[U_1|\boldsymbol{S}]} = R_1^2\sigma_1^2$. Then, notice that

$$\begin{aligned}
\mathrm{Cov}(U_1, \mathrm{E}[U_1 - U_0|\boldsymbol{S}]) &= \mathrm{Cov}(U_1, \mathrm{E}[U_1|\boldsymbol{S}]) - \mathrm{Cov}(U_1, \mathrm{E}[U_0|\boldsymbol{S}]) \\
&= \sigma_1^2 R_1^2 - \mathrm{Cov}(U_1, E[U_0|\boldsymbol{S}])
\end{aligned}$$

where the second equality follows from the fact that $\mathrm{Cov}(U_1, \mathrm{E}[U_1|\boldsymbol{S}]) = \mathrm{Var}(\mathrm{E}[U_1|\boldsymbol{S}]) = \sigma_1^2 R_1^2$. We therefore can identify $\mathrm{Cov}(U_1, E[U_0|\boldsymbol{S}])$. Next, we show that $\mathrm{Cov}(E[U_1|\boldsymbol{S}], E[U_0|\boldsymbol{S}])$ is identified. Notice that agents in the model utilize identical signals in predicting both U_1 and U_0 and that within the Gaussian model, the posterior mean is a linear function of the signals, which leads us to the following:

$$\mathrm{Cov}(U_1, \mathrm{E}[U_0|\boldsymbol{S}]) = \mathrm{Cov}(\mathrm{E}[U_1|\boldsymbol{S}] + \nu, \mathrm{E}[U_0|\boldsymbol{S}]) = \mathrm{Cov}(\mathrm{E}[U_1|\boldsymbol{S}], \mathrm{E}[U_0|\boldsymbol{S}]),$$

where ν is the the residual from projecting U_1 on $\boldsymbol{S}$ and satisfy $\mathrm{Cov}(\nu, \boldsymbol{S}) = 0$ Next, to identify R_0^2, notice that we can write the identified beliefs variance, σ_E^2, as:

$$\begin{aligned}
\sigma^2_{\mathrm{E}} &= \mathrm{Var}(E[U_1 - U_0|\boldsymbol{S}]) \\
&= \mathrm{Var}(E[U_1|\boldsymbol{S}]) + \mathrm{Var}(E[U_0|\boldsymbol{S}]) - 2\mathrm{Cov}(E[U_1|\boldsymbol{S}], E[U_0|\boldsymbol{S}]) \\
&= \sigma_1^2 R_1^2 + \sigma_0^2 R_0^2 - 2\mathrm{Cov}(E[U_1|\boldsymbol{S}], E[U_0|\boldsymbol{S}]).
\end{aligned}$$

which also allows us to identify $\sigma^2_{E[U_0|\boldsymbol{S}]}$. Next, using an equivalent argument to the identification of $\mathrm{Cov}(U_1, E[U_0|\boldsymbol{S})$, and R_0^2, we identify $\mathrm{Cov}(U_0, E[U_1|\boldsymbol{S}])$. Finally to identify $\mathrm{Cov}(U_d, \mathrm{E}[U_d|\boldsymbol{S}])$ we have the following equality:

$$\text{Cov}(U_d, \text{E}[U_d|\boldsymbol{S}]) = \text{Var}(E[U_d|\boldsymbol{S}]) = \sigma_d^2 R_d^2.$$

which concludes the proof. □

We now proceed to prove proposition 1.

Proof. Fix ρ, and let $C_S(\rho)$ be the implied covariance matrix of the signal vector $\boldsymbol{S}$ and potential earnings U_1 and U_0

$$C_S(\rho) = \begin{bmatrix} \sigma_{S_1,S_1} & \cdots & \sigma_{S_1,S_n} & \sigma_{S_1,U_1} & \sigma_{S_1,U_0} \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ \sigma_{S_n,S_1} & \cdots & \sigma_{S_n,S_n} & \sigma_{S_n,U_1} & \sigma_{S_n,U_0} \\ \sigma_{U_1,S_1} & \cdots & \sigma_{U_1,S_n} & \sigma_1^2 & \rho\sigma_0\sigma_1 \\ \sigma_{U_0,S_1} & \cdots & \sigma_{U_0,S_n} & \rho\sigma_0\sigma_1 & \sigma_0^2 \end{bmatrix},$$

and let $C_E(\rho)$ be the covariance matrix between marginal beliefs, $\text{E}[U_1|s]$, $\text{E}[U_0|s]$ and potential earnings, U_1 and U_0

$$C_E(\rho) = \begin{bmatrix} \sigma_{\text{E}_1}^2 & \sigma_{\text{E}_1,\text{E}_0} & \sigma_{\text{E}_1,U_1} & \sigma_{\text{E}_1,U_0} \\ \sigma_{\text{E}_0,\text{E}_1} & \sigma_{\text{E}_0}^1 & \sigma_{\text{E}_0,U_1} & \sigma_{\text{E}_0,U_0} \\ \sigma_{U_1,\text{E}_1} & \sigma_{U_1,\text{E}_0} & \sigma_1^2 & \rho\sigma_1\sigma_0 \\ \sigma_{U_0,\text{E}_1} & \sigma_{U_0,\text{E}_0} & \rho\sigma_1\sigma_0 & \sigma_0^2 \end{bmatrix}.$$

We next demonstrate that $C_S(\rho)$ is positive semi-definite (PSD) if and only if $C_E(\rho)$ is PSD. Without loss of generality, we focus on scenarios where signals are independent and possess unit variance. This approach is without loss, as for any feasible ρ, we can always residualize and rescale the signals, thereby maintaining their information content unchanged. We start by showing that if $C_E(\rho)$ is PSD, then matrix $C_S(\rho)$ is also PSD. We consider the contrapositive case and show that if matrix $C_S(\rho)$ is not PSD, then $C_E(\rho)$ is not PSD. Assume $C_S(\rho)$ not PSD. Then, there exists a vector t such that $t'C_S(\rho)t < 0$. Denote t_{s_i} the value in vector t that corresponds to signal s_i. and by t_1 and t_0 the value in vector t that correspond to U_1 and U_0. Using the fact that signals are uncorrelated, we can write

$$t'C_S(\rho)t = \sum_i t_{s_i}^2 + t_1 \left(\sum \sigma_{s_i,1} t_{s_i}\right) + t_0 \left(\sum \sigma_{s_i,0} t_{s_i}\right) + \sigma_1^2 t_1^2 + \sigma_0^2 t_0^2 + 2\rho\sigma_0\sigma_1 t_1 t_0 < 0 \tag{1.8}$$

We now show that there must exists a vector k, such that $k'C_E(\rho)k < 0$. Denote k_{E_d}, k_1 and k_0, similar to before, then

$$\begin{aligned} k'C_E(\rho)k &= 2k_1(\sigma_{\mathrm{E}_1}^2 k_{\mathrm{E}_1} + \sigma_{\mathrm{E}_1,\mathrm{E}_0} k_{\mathrm{E}_0}) \\ &+ 2k_0(\sigma_{\mathrm{E}_0}^2 k_{\mathrm{E}_0} + \sigma_{\mathrm{E}_1,\mathrm{E}_0} k_{\mathrm{E}_1}) \\ &+ (2\sigma_{1,0} k_{\mathrm{E}_1} k_{\mathrm{E}_0} + \sigma_{\mathrm{E}_1} k_{\mathrm{E}_1}^2 + \sigma_{\mathrm{E}_0} k_{\mathrm{E}_0}^2) \\ &+ \sigma_1^2 k_1^2 + \sigma_0^2 k_0^2 + 2\rho\sigma_1\sigma_0 k_1 k_0 \end{aligned}$$

As we restricted attention to the case where signals are uncorrelated and unit variance, and the conditional distribution of Gaussian model is linear function of signals, we can rewrite these expressions as

$$
\begin{aligned}
k'C_E(\rho)k &= 2k_1(k_{\mathrm{E}_1}\sum_{s_i}\sigma^2_{s_i,1} + k_{\mathrm{E}_0}\sum_{s_i}\sigma_{s_i,1}\sigma_{s_i,0}) \\
&+ 2k_0(k_{\mathrm{E}_0}\sum_{s_i}\sigma^2_{s_i,0} + k_{\mathrm{E}_1}\sum_{s_i}\sigma_{s_i,1}\sigma_{s_i,0}) \\
&+ (2k_{\mathrm{E}_1}k_{\mathrm{E}_0}\sum_{s_i}\sigma_{s_i,1}\sigma_{s_i,0}) + k^2_{\mathrm{E}_1}\sum_{s_i}\sigma^2_{s_i,1} + k^2_{\mathrm{E}_0}\sum_{s_i}\sigma^2_{s_i,0}) \\
&+ \sigma_1^2k_1^2 + \sigma_0^2k_0^2 + 2\rho\sigma_1\sigma_0k_1k_0 \\
&= 2k_1(\sum_{s_i}\sigma_{s_i,1}(\sigma_{s_i,1}k_{\mathrm{E}_1} + \sigma_{s_i,0}k_{\mathrm{E}_0})) \\
&+ 2k_0(\sum_{s_i}\sigma_{s_i,0}(\sigma_{s_i,0}k_{\mathrm{E}_0} + \sigma_{s_i,1}k_{\mathrm{E}_1})) \\
&+ \sum_{s_i}(\sigma_{s_i,1}k_{\mathrm{E}_1} + \sigma_{s_i,0}k_{\mathrm{E}_0})^2 \\
&+ \sigma_1^2k_1^2 + \sigma_0^2k_0^2 + 2\rho\sigma_1\sigma_0k_1k_0
\end{aligned}
$$

We now show how to find values of the vector k that makes this expression negative. We set $k_1 = t_1$ and $k_0 = t_0$. We use the additional two values of k to equate the remaining values such that $k'C_E(\rho)k = t'C_S(\rho)t < 0$. To do so, we notice we have two equation for two parameters

$$2\sum_{s_i}(\sigma_{s_i,1}k_{\mathrm{E}_1} + \sigma_{s_i,0}k_{\mathrm{E}_0}))(k_1\sigma_{s_i,1} + k_0\sigma_{s_i,0}) = \sum_{s_i}t_{s_i}(k_1\sigma_{s_i,1} + k_0\sigma_{s_i,0}) \tag{1.9}$$

and

$$\sum_{s_i}(\sigma_{s_i,1}k_{\mathrm{E}_1} + \sigma_{s_i,0}k_{\mathrm{E}_0})^2 = \sum_i t^2_{s_i} \tag{1.10}$$

Using the first equation, we can then solve for k_{E_1} in terms of known values and k_{E_0}

$$k_{\mathrm{E}_1} = \frac{\frac{1}{2}\sum_{s_i} t_{s_i}(k_1\sigma_{s_i,1} + k_0\sigma_{s_i,0}) - k_{\mathrm{E}_0}(k_1\sigma_{s_i,1} + k_0\sigma_{s_i,0})}{\sum_{s_i}\sigma_{s_i,1}}$$

plug this back into equation 1.10, we see that we have continuous function of k_{E_0}. This function goes from from 0 to infinity, the right hand side is a finite and positive expression, then by the Intermediate value theorem there exists a solution, which implies that there exists a vector for which $k'Bk < 0$ and B is not PSD. To show the reverse, we can follow the steps in reverse, and show that that if $C_E(\rho)$ is not PSD then $C_S(\rho)$ is not PSD as well, which concludes the proof. □

1.10.2 Nonparametric Identification of the Choice Model

We explore the non-parametric identification of choices. First, we identify the distribution of structural components, α_1 and α_0, by leveraging panel data, an instrumental variable, and specific wage structure assumptions. Next, we establish the identification of both the cost function and the beliefs distribution. While panel data aids in identifying α_1 and α_0. This step can be skipped if one assumes that outcomes are observed without measurement error.

In our analysis, we work under the assumption that the researcher has access to a random, independently and identically distributed sample of observations, each denoted by $(Y_{a,i}, D_i, X_i, Z_i)$. All analyses are conditional on the covariates vector X, so we omit the X notation for simplicity.

1.10.2.1 Identification of $P(\alpha_1, \mathrm{E}[\alpha_1 - \alpha_0|\boldsymbol{S}])$, $P(\alpha_0, \mathrm{E}[\alpha_1 - \alpha_0|\boldsymbol{S}])$ and the Threshold Function

We impose the following assumptions on the wage data generating Process. Wages are set according to

$$Y_{i,a} = D_i(\alpha_1 + \epsilon^1_{i,a}) + (1 - D_i)(\alpha_0 + \epsilon^0_{i,a}))$$

where $Y_{i,a}$ is individual i's income at age a, D_i is a dummy variable indicating whether the HG i attended four years college or not. One can think of α_d as individual fixed effect, if that individual goes to college or not. We further impose the following assumptions on the wage process

Assumption 1. *(1) for all a we have $E[\epsilon^D_{i,a}|\alpha_D] = 0$ (2) $\alpha_1, \alpha_0 \perp\!\!\!\perp \epsilon^D_{i,a}$ and (3) there exist at least two periods a^D, a'^D for each $D \in \{0, 1\}$ such that $\epsilon^D_{i,a} \perp\!\!\!\perp \epsilon^D_{i,a'}$*

Denote by $P(Z) = E[D = 1|Z]$ the propensity score conditional on Z. We then employ the following assumption

Assumption 2. *The characteristic functions of the conditional distribution $\alpha_1|D = 1, P(Z) = p$, $\alpha_0|D = 1, P(Z) = p$, $\epsilon^D_{i,a}|D = 1, P(Z) = p$ and $\epsilon^D_{i,a'}|D = 1, P(Z) = p$ are non vanishing*

The first part of Assumption 1 is standard and implies that any constant is absorbed into α^D, ensuring that deviations from the structural component are independent of the fixed effects. The second restriction mandates the existence of at least two periods in which the shocks are mutually independent, given the covariates X. While this condition is restrictive, it accommodates complex correlation structures, such as finite moving averages or other forms of multi-period correlations. The Assumption 2 stipulates that the characteristic functions of the conditional distributions for $\alpha_1|D = 1, P(Z) = p$, $\alpha_0|D = 1, P(Z) = p$, $\epsilon^D_{i,a}|D = 1, P(Z) = p$, and $\epsilon^D_{i,a'}|D = 1, P(Z) = p$ are non-vanishing[12]. This is a standard assumption that is used for nonparametric identification of factor models and assures us that we can use the characteristic functions to back-out the distribution of α_d.

Next, we impose restrictions on the agent information set. In the spirit of rational expectations, we assume that there are two parts to wages; a structural component, on which individuals have information on, and an unpredictable shock component that is not known to the high school gradutes.

Assumption 3 (Information Restriction)**.** *The signals individuals obtain do not contain any information on the non structural part of the wage,* $\epsilon^1_{i,a}, \epsilon^0_0$.

$$\boldsymbol{S}_i \perp\!\!\!\perp \epsilon^1_{i,a}, \epsilon^0_{i,a} | \alpha_1, \alpha_0$$

This implies that individuals can only receive information on the structural component of the wage, but may not have information on time varying shocks. Finally we impose the following assumptions on the instrument Z

Assumption 4 (Instrument Restrictions)**.** *We assume that the instrument satisfies the following conditions*

1. $\epsilon^1_{i,a}, \epsilon^0_{i,a}, \alpha_1, \alpha_0 \perp\!\!\!\perp Z$
2. $\boldsymbol{S} \perp\!\!\!\perp Z | \alpha_1, \alpha_0$
3. Z *is continuously distributed on* $\mathcal{Z} \subseteq \mathbb{R}$
4. $E[\alpha_1 - \alpha_0 | s]$ *continuously distributed*
5. $c(Z)$ *is differentiable with respect to* z *and covers the entire support of* $E[\alpha_1 - \alpha_0 | \boldsymbol{S}]$

12. The non vanishing assumption can be further relaxed, as shown in Evdokimov and White (2012)

The assumptions are akin to standard Instrumental Variable (IV) assumptions (Heckman and Vytlacil (2005)), but they incorporate additional structure through the modeling of the choice equation. The first assumption establishes the instrument's independence from the outcome variables. The second dictates that information is independent of the instrument, conditioned on the structural components. Notably, these first two assumptions collectively imply that $Y_1, Y_0, E[\alpha_1(t_i) - \alpha_0(t_i)|\boldsymbol{S}] \perp\!\!\!\perp Z$, aligning with standard IV assumptions where the selection variable is uncorrelated with the instrument. The final part of Assumption 4 is a technical requirement ensuring that we can recover the cost function by monitoring the derivative, as demonstrated in the proof.

Denote $E[\alpha_1 - \alpha_0|\boldsymbol{S}] = \mathrm{E}$. We show the following proposition.

Proposition 2. Under assumptions (1)-(4), $P(\alpha_1, \mathrm{E})$,$P(\alpha_0, \mathrm{E})$ and the cost function $c(z)$ are identified

Proof. Let a, a' be two periods such that $\epsilon^D_{i,a} \perp\!\!\!\perp \epsilon^D_{i,a}$. We start by showing how to identify $\mathrm{P}(\alpha_d|\mathrm{E})$ First, using assumption 1, 3, and 4 we have that $\epsilon^D_{i,a} \perp\!\!\!\perp \alpha_1|p(Z) = p, D = 1$ as

$$
\begin{aligned}
\mathrm{P}(\epsilon^D_{i,a}, \alpha_1|p(Z) = p, D = 1) &= \mathrm{P}(\epsilon^D_{i,a}, \alpha_1|p(Z) = p, \mathrm{E} \geq c(z)) \\
&= \mathrm{P}(\epsilon^D_{i,a}|\alpha_1, p(Z) = p, \mathrm{E} \geq c(z))\mathrm{P}(\alpha_1|p(Z) = p, \mathrm{E} \geq c(z)) \\
&= \mathrm{P}(\epsilon^D_{i,a}|p(Z) = p, \mathrm{E} \geq c(z))\mathrm{P}(\alpha_1|p(Z) = p, \mathrm{E} \geq c(z)) \\
&= \mathrm{P}(\epsilon^D_{i,a}|p(Z) = p, D = 1)\mathrm{P}(\alpha_1|p(Z) = p, D = 1)
\end{aligned}
$$

where the first equality stems from the choice model, the second stems from Bayes rule, and the third equality is due to the contraction rule and the decomposition rule of conditional Independence. We have an equivalent result for α_0 and $\epsilon_{i,a'}$. Last, notice that as $\epsilon_{i,a} \perp\!\!\!\perp \epsilon_{i,a'}$, $\epsilon_{i,a}, \epsilon_{i,a'} \perp\!\!\!\perp m_{\mathcal{R}}(s)$ and $\epsilon_{i,a}, \epsilon_{i,a'} \perp\!\!\!\perp Z$ we have that $\epsilon_{i,a} \perp\!\!\!\perp \epsilon_{i,a'}|p(Z) = p, D = 1$

Therefore $\epsilon^D_{i,a}$ and $\epsilon^D_{i,a'}$ and α_D are mutually independent, conditional on D and P, and we can now utilize Kotlarski's Lemma (1967) to identify the conditional distribution of α_1 and α_0. We first show how to identify the conditional distribution of α_1. Let $\Psi(y_a, y_{a'})$ be the conditional characteristic function of $(Y_{i,a}, Y_{i,a'})$ given $(P(z) = p, D = d)$. Let $\Psi_{\alpha_1}(t), \Psi_{\epsilon_a}(t)$ and $\Psi_{\epsilon'_a}(t)$ be the conditional characteristic functions of $\alpha_1, \epsilon_{i,a}, \epsilon_{i,a'}$, given $(P(z) = p, D = d)$, then we can show that (Rao (1992), page 21 and Gilraine et al. (2020))

$$log\Psi_{\alpha_1}(t) = iE[\alpha_1|D = 1, P(Z) = p]t + \int_0^t \frac{\partial}{\partial y_a}\left(log\frac{\Psi(y_a, y_{a'})}{\Psi(y_a, 0)\Psi(0, y_{a'})}\right)_{y_a=0} dy_{a'}$$

Noticing that

$$\frac{\partial}{\partial y_a}\left(log\frac{\Psi(y_a, y_{a'})}{\Psi(y_a, 0)\Psi(0, y_{a'})}\right)_{y_a=0} = \frac{\frac{\partial\Psi(0,y_{a'})}{\partial y_a}}{\Psi(0, y_{a'})} - iE[Y_{i,a}|D = 1, P(Z) = p]t$$

and that by assumptions 1 and 3 we have $iE[Y_{i,a}|D = 1, P(Z) = p]t = iE[\alpha_1|D = 1, P(Z) = p]$ we then get

$$\log\Psi_{\alpha_1}(t) = \int_0^t \frac{\frac{\partial\Psi(0,y_{a'})}{\partial y_a}}{\Psi(0, y_{a'})} dy_{a'}$$

as the characteristic function fully defines the distribution and $\Psi(y_a, y_{a'})$ is observed in the data, we have identified $\mathrm{P}(\alpha_1|D = 1, P(z) = p)$. Similar argument shows that we can identify $\mathrm{P}(\alpha_0|D = 0, P(z) = p)$.

Next, denote by $F_{\alpha_1}(\cdot|D = 1, P(Z) = p)$ the conditional CDF of α_1. Denote by $V = F_{\mathrm{E}}(\mathrm{E})$ the quantile of the beliefs in the beliefs distribution. Then following the arguments in Carneiro and Lee

(2009) we have that for all k on the support of α_1 we have that

$$\begin{aligned} F_{\alpha_1}(k|P(z), D=1) = E[\mathbb{1}\{\alpha_1 \leq k\}|P(Z)=p, D=1] &= E[\mathbb{1}\{\alpha_1 \leq k\}|P(Z)=p, V>p(Z)] \\ &= \frac{1}{p}\int_{1-p}^{1} E[\mathbb{1}\{\alpha_1 \leq k\}|V=v]dv \end{aligned}$$

rewriting the equation gives us

$$pE[\mathbb{1}\{\alpha_1 \leq k\}|P(Z)=p, D=1] = \int_{1-p}^{1} E[\mathbb{1}\{\alpha_1 \leq k\}|V=v]f(v)dv$$

Using assumption 4 we can take derivative from both sides, with respect to p, and get

$$E[\mathbb{1}\{\alpha_1 \leq k\}|V=1-p] = E[\mathbb{1}\{\alpha_1 \leq k\}|P(Z)=p, D=1] + p\frac{E[\mathbb{1}\{\alpha_1 \leq k\}|P(Z)=p, D=1]}{\partial p}$$

Therefore we have that $\mathrm{P}(\alpha_1|V)$ is identified. Following similar steps we have that $\mathrm{P}(\alpha_0|V)$ is also identified

$$E[\mathbb{1}\{\alpha_0 \leq k\}|V=1-p] = E[\mathbb{1}\{\alpha_0 \leq k\}|P(Z)=p, D=0] - (1-p)\frac{E[\mathbb{1}\{\alpha_0 \leq k\}|P(Z)=p, D=0]}{\partial p}$$

Next, observe that we can construct the probabilities $P(\alpha_1|\mathrm{E})$ and $P(\alpha_0|\mathrm{E})$ using the law of iterated expectations we have

$$e = E[\alpha_1 - \alpha_0|\mathrm{E}=e] = E[\alpha_1 - \alpha_0|F_{\mathrm{E}}(\mathrm{E}) = V] = \int \alpha_1 P(\alpha_1|V)d\alpha_1 - \int \alpha_0 P(\alpha_0|V)d\alpha_0. \quad (1.11)$$

Therefore we can identify the inverse, $F_{\mathrm{E}}^{-1}(V)$, and consequently the CDF of beliefs, $F_{\mathrm{E}}(e)$. As the CDF is strictly increasing and therefore invertible, by assumption 4 we can also identify $P(\alpha_1|\mathrm{E})$

and $P(\alpha_0|\mathrm{E})$ as needed. Therefore we've identified the joint $P(\alpha_1, \mathrm{E})$ and $P(\alpha_0, \mathrm{E})$. Finally, to identify the cost function, observe that

$$P(z) = \Pr(\mathrm{E} > c(z)) = 1-F(c(z)) \implies F_{\mathrm{E}}^{-1}(1-P(z)) = F_{\mathrm{E}}^{-1}(F(c(z))) \implies F_{\mathrm{E}}^{-1}(1-P(z)) = c(z).$$

Finally, in order to identify the cost function, we notice that

$$P(z) = \mathrm{P}(\mathrm{E} > c(z)) \quad = 1 - F(c(z)) \implies F_{\mathrm{E}}^{-1}(1 - P(z)) = F_{\mathrm{E}}^{-1}(F(c(z))) \implies F_{\mathrm{E}}^{-1}(1 - P(z)) = c(z)$$

which concludes the proof. □

1.10.2.2 A Testable Implication

As discussed in Canay et al. (2020) and in Hull (2021), the choice model implies that the Marginal Treatment effect (Heckman and Vytlacil (2005)) estimated using the instruments, should be decreasing. To see that notice that we use 1.11 to identify the CDF of V, therefore, if we get that this is not increasing function of v, this implies that our model is mispecificed. In the Gaussian model we estimate in the text this amounts to requiring that

$$\sigma_{\mathrm{E}} = \gamma_c^1 - \gamma_c^0 \geq 0,$$

as σ_{E} standard error.

1.11 Estimation

We now turn to describe how we estimate the Gaussian choice model. We first start by estimating α_1 and α_0 by averaging wages over periods of time

$$\hat{\alpha_{di}} = \frac{1}{T-t}\sum_{a=t}^{T} Y_{i,a}^{d}$$

Then, given our $\hat{\alpha}_1$ and $\hat{\alpha}_0$, we estimate the model in three steps. In the first step, we estimate the propensity score using a Probit model, the covariates X, and the instrument Z. In the second step, we use the Heckman control function approach (Heckman (1979)) to estimate β_1 and β_0. As discussed in the previous section, we obtain the standard deviation of beliefs from the coefficients on the control function. Next, we show how we can estimate the cost function. Using equation 1.4, we see that the Probit regression coefficients, standardized by the standard deviation of beliefs, are impacted by both beliefs and costs. To adjust for this, we rescale the coefficients and add the conditional expectations, estimated using the control function approach:

$$\hat{c}(x,z) = \hat{\sigma}_\eta \times (z\hat{b}_z + x\hat{b}_x) + x(\hat{\beta}_1 - \hat{\beta}_0).$$

Finally, to get σ_1 and σ_0, we solve the maximum likelihood function as shown in equation 1.5.

To our measure of information contribution to the gap we simply calculate the $\hat{R}^2$, as discussed in 1.2.3 for both groups. We then estimate the information channel as

$$\underbrace{\hat{P}(D=1|b,x)}_{\text{Observed}} - \underbrace{\frac{1}{N}\sum_i \Phi\left(\frac{x(\hat{\beta}_1 - \hat{\beta}_0) - \hat{c}(x_i,z_i)}{\sqrt{\hat{\sigma}^2_{\mathcal{R}}\hat{R}^2_a}}\right)}_{\text{Counterfactual}},$$

where the first part is just the observed share and the second part is the counterfactual share of individuals who choose to attend, if they had the same R^2 as group a. To estimate this part we simply average over $\Phi\left(\frac{x(\hat{\beta}_1-\hat{\beta}_0)-\hat{c}(x_i,z_i)}{\sqrt{\hat{\sigma}^2_{\mathcal{R}}\hat{R}^2_a}}\right)$ for all the observation of group b. Estimation of the composition channel is done the same.

1.12 Another Measure for Information Differences - Equating Information Structure Across Groups

1.12.1 Decomposition - The Role of Differences in Information Structures

In the main text we considered two ways to measure the role of information frictions on choice. We now consider an additional one that aims to equate the information structure across groups. Information structure is a tuple $\mathcal{S} = (P(s|\mathcal{R}), S)$ containing a set of conditional density function, that describes the probability of observing signal s, for an individual with return $\mathcal{R}$, and the support of these signals S. Information structures are widely used in economics and captures the mapping between the the state variables and beliefs (Bergemann and Morris (2016),Bergemann and Morris (2019)). In the following exercise we want to understand how the fact that different groups have access to different information structures, affect the choice gap. We therefore consider equating the information structure across the two groups. We then preform similar decomposition exercise as we did in section 1.2.4. In this decomposition exercise we decompose the choice gap to differences in choice that are attributed to differences arising from differences in the information structure and differences in the returns distribution:

$$\begin{aligned}
&\underbrace{P(D=1|\text{Group } b) - P(D=1|\text{Group } a)}_{\text{Total Effect}} = \\
&\underbrace{P(D=1|\text{Group } b) - P(D=1|\text{Group } b \text{ with information structure of Group } a)}_{\text{Information Channel}} \\
+&\underbrace{P(D=1|\text{Group } b \text{ with information structure of Group } a) - P(D=1|\text{Group } a)}_{\text{Composition Channel}} = \\
&\underbrace{\int_{\mathcal{R}\times c} \mathcal{P}(E_{b,b}(s) \geq c|\mathcal{R}, c, b)\pi_b(\mathcal{R}, c)d\mathcal{R}dc - \int_{\mathcal{R}\times c} \mathcal{P}(E_{a,b}(s) \geq c|\mathcal{R}, c, a)\pi_b(\mathcal{R}, c)d\mathcal{R}dc}_{\text{Information Channel}} \\
&+\underbrace{\int_{\mathcal{R}\times c} \mathcal{P}(E_{a,b}(s) \geq c|\mathcal{R}, c, a)\pi_b(\mathcal{R}, c) - \int_{\mathcal{R}\times c} \mathcal{P}(E_{a,a}(s) \geq c|\mathcal{R}, a)\pi_a(\mathcal{R}, c)d\mathcal{R}dc,}_{\text{Composition Channel}}
\end{aligned}$$

where

$$E_{a,a}(s) = \int_{\mathcal{R}} \mathcal{R} \frac{P(s|\mathcal{R}, a) \times \pi_a(\mathcal{R})}{\int_{\tilde{\mathcal{R}}} P(s|\tilde{\mathcal{R}}, a) \times \pi_a(\tilde{\mathcal{R}})d\tilde{\mathcal{R}}} d\mathcal{R}$$

is simply the beliefs of group b, when they have access to information of group b and prior of group b,[13] and

$$E_{a,b}(s) = \int_{\tilde{\mathcal{R}}} \tilde{\mathcal{R}} \frac{\overbrace{P(s|\tilde{\mathcal{R}}, a)}^{\text{Information}} \times \overbrace{\pi_b(\tilde{\mathcal{R}})}^{\text{earnings}}}{\int P(s|\tilde{\mathcal{R}}, a)\pi_b(\tilde{\mathcal{R}})d\tilde{\mathcal{R}}} d\tilde{\mathcal{R}}$$

is a counterfactual beliefs for group of members b, is they have the information structure of group a, but returns distribution of group b.

The information channel measures the extent to which the gap in choices would change if both groups had access to the same information structure as group a. Disparities in information structure

13. Remember that we assume rational expectations, hence the prior is the true distribution of returns

can arise from various environmental factors affecting the decision-maker. For instance, if members of group b typically have more academically inclined parents than those in group a, they are likely to receive more accurate information about the benefits of college for an individual, thus providing clearer signals about potential earnings post-college. Additionally, if the social networks of group b members are closely connected to a specific industry that requires certain information, this could create differences in individuals' abilities to predict returns. Therefore, the information channel quantifies the extent of the gap in choices that is attributable to individuals in the two groups receiving different signals, despite having equal potential returns.

It's essential to note two things. First, the information structure captures not only 'measurement' type signals, of the form *Signal* = *True value* + *Measurement Error*, as commonly seen in the literature, but also incorporates more sophisticated cases, that incorporate what individuals know and understand about the data-generating process. Second, in our decomposition exercise, we impose that individuals update their beliefs correctly. They use the new signals and their correct priors to adjust their understanding. In other words, we examine how they would update their beliefs knowing that the distribution of signals they receive comes from a new source.

The following examples demonstrate two points. First, how information structure incorporates the underlying data generating process that govern the returns, and is not simply a "measurement error" type signals. Second, the example shows that what's important is not equating the access to signals, but equating the meaning that these signals have, captured by the information structure.

Example 1.12.1 (Occupation and Earnings)**.** The informational content of the signals individuals might be more dependent on the structure of the economy itself. For instance, consider the case where the earnings of non-college-goers are zero for both members of groups a and b, and there are two occupations in the economy: lawyers and accountants. Both lawyers and accountants are paid either a high or low wage, $H > 0 > L$, with equal probability. Prior to deciding to go to college, individuals receive an informative signal on their potential returns if they end up being lawyers. Denote these signals as $\tilde{H}_{\text{law}}$ and $\tilde{L}_{\text{law}}$. The distributions of occupations, earnings, and the signal for each group are given below.

Group b

		H	**L**
Lawyer	$\tilde{H}_{law}$	$\frac{6}{20} \times \frac{5}{6}$	$\frac{6}{20} \times \frac{1}{6}$
	$\tilde{L}_{law}$	$\frac{6}{20} \times \frac{1}{6}$	$\frac{6}{20} \times \frac{5}{6}$
Accountant	$\tilde{H}_{law}$	$\frac{4}{20} \times \frac{1}{2}$	$\frac{4}{20} \times \frac{1}{2}$
	$\tilde{L}_{law}$	$\frac{4}{20} \times \frac{1}{2}$	$\frac{4}{20} \times \frac{1}{2}$

Group a

		H	**L**
Lawyer	$\tilde{H}_{law}$	$\frac{4}{20} \times \frac{5}{6}$	$\frac{4}{20} \times \frac{1}{6}$
	$\tilde{L}_{law}$	$\frac{4}{20} \times \frac{1}{6}$	$\frac{4}{20} \times \frac{5}{6}$
Accountant	$\tilde{H}_{law}$	$\frac{6}{20} \times \frac{1}{2}$	$\frac{6}{20} \times \frac{1}{2}$
	$\tilde{L}_{law}$	$\frac{6}{20} \times \frac{1}{2}$	$\frac{6}{20} \times \frac{1}{2}$

Table 1.18: Demonstration of Information Structure

In this economy, the share of high earners and low earners is $\frac{1}{2}$ for both groups. The share of individuals in both groups with signals $\tilde{H}_{\text{law}}$ and $\tilde{L}_{\text{law}}$ is also $\frac{1}{2}$. Moreover, for both groups, individuals who end up as lawyers and received a high signal have a $\frac{5}{6}$ probability of having high earnings. The only difference between the two groups is the share of individuals who end up being lawyers, versus those ending up being accountants. This difference implies that the signals each individual from each group receives have different information content, generating differences in the distribution of beliefs. For members of group b, the information structure is given by

$$P(\tilde{H}|H) = P(\tilde{H}|\text{Lawyer}, H)P(\text{Lawyer}|H) + P(\tilde{H}|\text{Acc}, H)P(\text{Acc}|H) = \frac{7}{10} \tag{1.12}$$
$$P(\tilde{H}|L) = P(\tilde{H}|\text{Lawyer}, L)P(\text{Lawyer}|L) + P(\tilde{H}|\text{Acc}, L)P(\text{Acc}|L) = \frac{3}{10} \tag{1.13}$$

Similarly, for members from group a we have

$$P(\tilde{H}|H) = P(\tilde{H}|\text{Lawyer}, H)P(\text{Lawyer}|H) + P(\tilde{H}|\text{Acc}, H)P(\text{Acc}|H) = \frac{19}{30} \tag{1.14}$$
$$P(\tilde{H}|L) = P(\tilde{H}|\text{Lawyer}, L)P(\text{Lawyer}|L) + P(\tilde{H}|\text{Acc}, L)P(\text{Acc}|L) = \frac{11}{30} \tag{1.15}$$

which implies that even when the marginal distribution of the signal and returns is the same, the implied beliefs given the same signal are different

$$m_{R_{\langle b,b\rangle}}(\tilde{H}) = H \times \frac{7}{10} + L \times \frac{3}{10} \tag{1.16}$$
$$m_{R_{\langle a,a\rangle}}(\tilde{H}) = H \times \frac{19}{30} + L \times \frac{11}{30} \tag{1.17}$$

Therefore, although the marginal distribution of signals and returns is the same in the economy, the information structure is different, and the same signal would be interpreted differently in both cases. What does it mean to switch the information structure between group a and group b in this environment? In the thought experiment we perform here, we ask what would be the observed behavior if we provided a signal with the same informational content on the returns as the other group. In this sense, our decomposition approach is "reduced form" in spirit, as we do not describe what drives the differences in information. Instead, we explore the ways in which systemic differences in information on earnings are provided to individuals and how they affect the observed gaps in behavior.

These differences can arise from various channels, some due to the way the economy is structured, others might be due to differences in individuals, such as the ability to process information or the financial ability to acquire information.

The following example shows that the same component can play a role as both a piece of information and part of the data generating process of the outcomes.

Example 1.12.2 (Knowledge of some structural components)**.** In some cases, individuals may know specific parts of the data-generating process of earnings. For example, assume that the earnings are determined by a function with a known component to the decision-maker, x, and some unknown component ν_d:

$$\alpha^1 = m_1(x, \nu_1) \tag{1.18}$$

$$\alpha^0 = m_0(x, \nu_0) \tag{1.19}$$

Here, x could represent known ability, latent cost of effort, or parental connections in the labor market. In this case, the information structure is simply the probability of observing x, given the earnings $P(x|\alpha^1, \alpha^0)$. This assumption is common in economic models where we believe that some variables affecting the outcomes are known to the decision-makers while making choices, and they use them to form beliefs about the outcomes. Therefore, in our thought experiment of switching the information structure between groups, we separate the two roles of x. Specifically, we fix the distribution of x in the population, thereby keeping the distribution of earnings fixed. But we ask what would happen if the agent did not know x, but instead had access to a similar information environment as group a, and how that would change choice patterns.

We now proceed to explore the second component of decomposition - the composition channel.

We can express this channel as:

$$\underbrace{P_{\langle a|b\rangle} - P_{\langle a|a\rangle}}_{\text{Composition Channel}} = \int_{\alpha^1,\alpha^0} P(m_{R_{\langle a,b\rangle}}(s) \geq 0|\mathcal{R}, a)\frac{\pi(\mathcal{R}|b)}{\pi(\mathcal{R}|a)}\pi(\mathcal{R}|a) \tag{1.20}$$

$$- \int_{\mathcal{R}} P(m_{R_{\langle a,a\rangle}}(s) \geq 0|\mathcal{R}, a)\pi(\mathcal{R}|a)d\alpha^1 d\alpha^0 \tag{1.21}$$

In the composition channel, we maintain the information structure of group a, yet re-weight the population of group a to align with the distribution of group b. This thought experiment explores how the share of college attenders from group a would change if we modified the composition of the group, so that their distribution of earnings would align with that of group b. In this counterfactual, we are not breaking the connection between information and earnings, as we did in example 2.3, but merely shifting the proportion of individuals at certain earnings levels, ensuring that they take the change into account while forming their beliefs. As we alter the distribution of earnings, while keeping the information structure fixed, we also modify the marginal distribution of signals within the population. This means that if, for instance, we increased the proportion of potential students with high $\mathcal{R}$, we are also enlarging the population's share of those receiving signals tied to higher earnings. Consequently, maintaining the information structure fixed means that we are transforming the distribution of signals in the population, but keep it's meaning.

Example 1.12.3 (Knowledge of some structural components-Continued)**.** In this example, the composition channel involves adjusting the share of members in group a with specific earnings levels, to align with those from group b. It's important to note that we are not necessarily equalizing the share of variable x between the two groups. If x represents, for example, ability, and the function $m(., \nu)$ varies between groups, our hypothetical scenario doesn't balance the share of high and low ability across both groups. If m differs, matching the share of high and low ability could result in significantly different distributions. Since the HGs are not concerned with ability itself but as an indicator of their returns, aligning individual parts across groups doesn't provide insight into how the distribution of outcomes influences choice.

Similar to our discussion in the main text, it's crucial to understand that our analysis offers a partial equilibrium perspective on changing information structure. The information structure in many cases changes endogenously.For example, individuals may exert effort to generate better information in response to the distribution of returns. It also may be that differences in information could arise due to selection and equilibrium effects. For example, if information influences labor market selection patterns, and employers respond to these patterns, our counterfactuals won't address this. Our analysis assumes that the existing information structure is a given and demonstrates further details in the appendix.

1.12.1.1 Gaussian Scalar Interpretation

In the scalar Gaussian case we can write the decomposition explicitly as

$$P(D=1|b) - P(D=1|a) = \underbrace{\int_X \Phi\left(\frac{\mu_{\mathcal{R},b,x} - c_b(x)}{\sqrt{\frac{\sigma^4_{\mathcal{R},b,x}}{\sigma^2_{\mathcal{R},b,x}+\sigma^2_{\epsilon,b,x}}}}\right) dF_b(x) - \int_X \Phi\left(\frac{\mu_{\mathcal{R},b,x} - c_b(x)}{\sqrt{\frac{\sigma^4_{\mathcal{R},b,x}}{\sigma^2_{\mathcal{R},b,x}+\sigma^2_{\epsilon,a,x}}}}\right) dF_b(x)}_{\text{Information Channel}}$$

$$+ \underbrace{\int_X \Phi\left(\frac{\mu_{\mathcal{R},b,x} - c_b(x)}{\sqrt{\frac{\sigma^4_{\mathcal{R},b,x}}{\sigma^2_{\mathcal{R},b,x}+\sigma^2_{\epsilon,a,x}}}}\right) dF_b(x) - \int_c \Phi\left(\frac{\mu_{\mathcal{R},a,x} - c_a(x)}{\sqrt{\frac{\sigma^4_{\mathcal{R},a,x}}{\sigma^2_{\mathcal{R},a,x}+\sigma^2_{\epsilon,a,x}}}}\right) dF_a(x)}_{\text{Composition Channel}}$$

Therefore, in the scalar Gaussian case, equating information structure across two groups essentially means equalizing the level of uncertainty surrounding true returns.

Remark. Notice that in our discussion here we fixed the information structure, as signals conditional on returns. We did this, as returns are what agents care about, and for the decision process they are indifferent between two pairs of earnings with the same difference. Therefore from the perspective of the agents, the payoff relevant value for the decision is the difference. Another approach can be to define the information structure on earnings. This would imply a different interpretation of information.

In the following parts we discuss how this decomposition measure can identified under different assumption on the data or the type of fundamentals and information.

1.12.2 Nonparametric point Identification of the Decomposition Components

Fix two groups $g \in \{a, b\}$. In the subsequent sections, we demonstrate how to identify the quantity

$$P_{\langle a,b\rangle} = \int_{\mathcal{R}\times c} \mathcal{P}(E_{a,b}(s) \geq c|\mathcal{R}, c, a)\pi_b(\mathcal{R}, c) \tag{1.22}$$

required for decomposition. As outlined in Section 1.10.2, the primary challenge lies in constructing the distribution of posterior means that incorporates both the counterfactual distribution of signals and returns. This must be achieved despite having access only to the conditional expectations distribution, rather than the complete information structure available to agents. We first establish conditions for point identification, then extend our analysis to more general cases for identifying this quantity. Throughout the analysis we assume that $\pi(c)$ is identified, and implicitly condition on the cost.

1.12.2.1 Point Identification Under Increasing Beliefs Function

We start by showing that if we are willing to assume that the information is scalar, and that beliefs are increasing function of that signal, then the quantity in 1.22 is identified.

Proposition 3. Let $E[\mathcal{R}|s]$ be a strictly increasing function of s, then equation 3 is identified.

Proof. The claim follows trivially from the fact that a strictly monotonic transformation is merely a renaming of the signal but does not alter its information content. Therefore, individuals update beliefs in the same manner, using either the information structure's likelihood functions $P(s|\mathcal{R})$ with support $\mathcal{S}$ or $P(E[\mathcal{R}|s]|\mathcal{R})$ with support given by the posterior means, for any prior. To illustrate this in our continuous density of signals case, we have

$$E_{a,b}(s) = \int \mathcal{R} \frac{P_a(s|\mathcal{R})\pi_b(\mathcal{R})}{\int P_a(s|\mathcal{R})\pi_b(\mathcal{R})d\mathcal{R}} d\mathcal{R} = \int \mathcal{R} \frac{\left|\frac{1}{\frac{\partial E_a}{\partial s}}\right| P_a(E_a|\mathcal{R})\pi_b(\mathcal{R})}{\left|\frac{1}{\frac{\partial E_a}{\partial s}}\right| \int P_a(E_a|\mathcal{R})\pi_b(\mathcal{R})d\mathcal{R}} d\mathcal{R} = E[\mathcal{R}|E_a(s); b]$$

where E_a denotes the beliefs of group a, with their information structure and prior, and $E[\mathcal{R}|E_a(s); b]$ is the belief induced by observing the signal $E_a(s)$ and prior π_b. As demonstrated in section 1.10.2, for a given correlation between α_1 and α_0, we can identify the joint distribution of $E[\mathcal{R}|s]$ and $\mathcal{R}$ for groups a and b. Therefore, as each signal corresponds to a unique belief, we can calculate the implied counterfactual beliefs distribution directly from the identified distribution of beliefs. Consequently, $P(E_{a,b}(s)|\mathcal{R})$ is identified, and equation 1.22 is trivially identified. $\square$

Under what conditions can we expect the conditional expectation to be a strictly increasing function of returns? A sufficient condition for this is that the joint distribution of $\mathcal{R}$ and s satisfies the Monotone Likelihood Ratio Property (MLRP). The following corollary formalizes this claim. Let $P(\mathcal{R}, s)$ satisfy the strict Monotone Likelihood Ratio Property,

$$\forall s > s', x > x' \quad P(\mathcal{R}|s)P(\mathcal{R}'|s') > P(\mathcal{R}'|s)P(\mathcal{R}|s') \tag{1.23}$$

then the quantity in equation 3 is identified.

Proof. The corollary follows from the preceding proposition and the fact that MLRP implies First-Order Stochastic Dominance,

$$F_s(\mathcal{R}) \leq F_{s'}(\mathcal{R}) \tag{1.24}$$

which implies that the conditional expectation is strictly increasing,

$$E[\mathcal{R}|s] = \int_{\mathcal{R}} (1 - F_s(\mathcal{R}))\, d\mathcal{R} > \int_{\mathcal{R}} (1 - F_{s'}(\mathcal{R}))\, d\mathcal{R} = E[\mathcal{R}|s'].$$

Here, F_s denotes the CDF of $\mathcal{R}$, conditional on s. $\square$

1.12.2.2 Identification Under the General Gaussian Model

We reintroduce the Gaussian model as in section 1.2. Throughout the discussion, we fix the cost c and make the identification argument conditional on c. Again, we assume that individuals observe a scalar signal S, and the structural components of earnings α_1, α_0 are drawn from a joint normal distribution

$$\begin{pmatrix} \boldsymbol{S} \\ \alpha_1 \\ \alpha_0 \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \boldsymbol{\mu}_s \\ \mu_1 \\ \mu_0 \end{pmatrix}, \begin{bmatrix} \Sigma_{\boldsymbol{S}}, \Sigma_{\boldsymbol{S},1}, \Sigma_{\boldsymbol{S},0} \\ \Sigma_{\boldsymbol{S},1}, \sigma_1, \sigma_{1,0} \\ \Sigma_{\boldsymbol{S},0}, \sigma_{1,0}, \sigma_0 \end{bmatrix}\right)$$

Using the properties of the normal distribution, we can write the joint distribution of the signals and the returns, where $\mathcal{R} = \alpha_1 - \alpha_0$, as

$$\begin{pmatrix} \boldsymbol{S} \\ \mathcal{R} \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \boldsymbol{\mu}_s \\ \mu_1 - \mu_0 \end{pmatrix}, \begin{bmatrix} \Sigma_{\boldsymbol{S}} & \Sigma_{\boldsymbol{S},\mathcal{R}} \\ \Sigma_{\boldsymbol{S},\mathcal{R}}^T & \sigma_1^2 + \sigma_0^2 - 2\sigma_{1,0} \end{bmatrix}\right)$$

Where $\Sigma_{\boldsymbol{S},\mathcal{R}} = \Sigma_{\boldsymbol{S},1} - \Sigma_{\boldsymbol{S},0}$. Given a signal realization $\boldsymbol{S}$, the information structure, $\Pr(\boldsymbol{S}|\mathcal{R})$, is then given by

$$\Pr(\boldsymbol{S}|\mathcal{R}) = \mathcal{N}\left(\mu_{\boldsymbol{S}} + \Sigma_{S,\mathcal{R}}\sigma_{\mathcal{R}}^{-2}(\mathcal{R} - \mu_{\mathcal{R}}), \Sigma_{\boldsymbol{S}} - \Sigma_{\boldsymbol{S},\mathcal{R}}\sigma_{\mathcal{R}}^{-2}\Sigma_{\boldsymbol{S},\mathcal{R}}^T\right)$$

An individual with signal realization $\boldsymbol{S}$ forms the following posterior mean:

$$E[\mathcal{R}|\boldsymbol{S}] = \mu_{\mathcal{R}} + \Sigma_{S,R}^{T}\Sigma_{\boldsymbol{S}}^{-1}(\boldsymbol{S} - \mu_{S})$$

This implies that individuals i with cost c and signal realization $\boldsymbol{S}$ would choose to go to college if

$$D = \mathbb{1}\left[E[\alpha_1 - \alpha_0|\boldsymbol{S}] \geq c\right] = \mathbb{1}\left[\mu_{\mathcal{R}} + \Sigma_{S,R}^{T}\Sigma_{\boldsymbol{S}}^{-1}(\boldsymbol{S} - \mu_{S}) \geq c\right]$$

We can calculate the share of students who attend college with cost c. First, we note that the beliefs distribution is given by

$$E[\mathcal{R}|\boldsymbol{S}] \sim \mathcal{N}\left(\mu_{\mathcal{R}}, \Sigma_{S,\mathcal{R}}^{T}\Sigma_{S}^{-1}\Sigma_{S,\mathcal{R}}\right)$$

Therefore, the share of individuals who would go to college is given by

$$P(D = 1|c) = \Phi\left(\frac{\mu_{\mathcal{R}} - c}{\Sigma_{S,\mathcal{R}}^{T}\Sigma_{S}^{-1}\Sigma_{S,\mathcal{R}}}\right)$$

Now, again, we assume that individuals are divided into two groups $g \in \{a, b\}$. Fixing a copula parameter between α_1 and α_0 for each group, and using results from section 1.10.2, we know we can identify the joint distribution of returns and beliefs for groups a and b, $P_a(\mathcal{R}, E(s))$ and $P_b(\mathcal{R}, E(s))$. We now show that this is sufficient to identify the quantity in 3 and solve for the decomposition.

Given the information structure of group a, we can derive the counterfactual joint distribution of signals and returns as follows[14]

$$\begin{pmatrix} \boldsymbol{S}_a \\ \mathcal{R}_b \end{pmatrix} \sim N\left(\begin{pmatrix} k_a + m_a\mu_{\boldsymbol{S}_a} \\ \mu_{\mathcal{R}_b} \end{pmatrix}, \begin{bmatrix} m_a\sigma_b^2 {m_a}^T + \Sigma_{\boldsymbol{S}_a} - m_a\Sigma_{\boldsymbol{S}_a,\mathcal{R}_b}^T & m_a\sigma_{\mathcal{R}_b}^2 \\ {m_a}^T\sigma_{\mathcal{R}_b}^2 & \sigma_{\mathcal{R}_b}^2 \end{bmatrix}\right)$$

where $k_a = \mu_{S_a} - \Sigma_{\boldsymbol{S}_a,\mathcal{R}_b}\sigma_{\mathcal{R}_b}^{-2}\mu_{\mathcal{R}_b}$ and $m_a = \Sigma_{\boldsymbol{S}_a,\mathcal{R}_b}\sigma_{\mathcal{R}_b}^{-2}$ and subscript $g \in \{a, b\}$ indicates that the parameters are from the distribution of group g.

We can now derive the counterfactual posterior mean belief, given a signal realization $\boldsymbol{S}$.

$$E_{a,b} = \mu_b + m_a^T\sigma_{\mathcal{R}_b}^2\left(\Sigma_{\boldsymbol{S}_a,\mathcal{R}_a}\sigma_{\mathcal{R}_a}^{-2}\sigma_{\mathcal{R}_b}^2\sigma_{\mathcal{R}_a}^{-2}\Sigma_{\boldsymbol{S}_a,\mathcal{R}_a}^T + \Sigma_{\boldsymbol{S}_a} - \Sigma_{\boldsymbol{S}_a,\mathcal{R}_a}\sigma_{\mathcal{R}_a}^{-2}\Sigma_{\boldsymbol{S}_a,\mathcal{R}_a}^T\right)^{-1}\left(\boldsymbol{S}_a - k_a - m_a\mu_{\boldsymbol{S}_a}\right)$$

and the counterfactual belief distribution is given by

$$E_{a,b} \sim N\left(\mu_b, \sigma_{\mathcal{R}_b}^4 m_a^T\left(\left(m_a\sigma_b^2 m_a^T + \Sigma_{\boldsymbol{S}_a} - m_a\Sigma_{\boldsymbol{S}_a,\mathcal{R}_a}^T\right)^{-1}\right)^T m_a\right)$$

Denote by OV_a the identified variance of beliefs for group a

$$OV_a = \Sigma_{S_a,\mathcal{R}_a}^T\Sigma_{S_a}^{-1}\Sigma_{S_a,\mathcal{R}_a}$$

The following proposition assert that we can identify the variance of the counterfactaul beliefs distribution

Proposition 4. Let $\mathcal{R}$ and signal vector $\boldsymbol{S}$ be jointly Gaussian-distributed, conditional on the cost

14. We slightly abuse notation here setting $\mathcal{R}_b$ to denote that returns are distributed as in group b

c, for members of both group a and b. Then we we can point identify the counterfactual quantity

$$\int_{\mathcal{R}\times c} \mathcal{P}(E_{a,b}(s) \geq c|\mathcal{R}, c, a)\pi_b(\mathcal{R}|c)p(c)d\mathcal{R}dc$$

Proof. The proof follows from the following derivation:

$$\begin{aligned}
\mathrm{Var}(E_{a,b}) &= \sigma^4_{\mathcal{R}_b} m_a^T \left(\left(\Sigma_{\boldsymbol{S}_a,\mathcal{R}_a} \sigma^{-2}_{\mathcal{R}_a} \sigma^2_{\mathcal{R}_b} \sigma^{-2}_{\mathcal{R}_a} \Sigma^T_{\boldsymbol{S}_a,\mathcal{R}_a} + \Sigma_{\boldsymbol{S}_a} - \Sigma_{\boldsymbol{S}_a,\mathcal{R}_a} \sigma^{-2}_{\mathcal{R}_a} \Sigma^T_{\boldsymbol{S}_a,\mathcal{R}_a} \right)^{-1} \right)^T m_a \\
&= \frac{\sigma^4_{\mathcal{R}_b}}{\sigma^4_{\mathcal{R}_a}} \Sigma^T_{\boldsymbol{S}_a,\mathcal{R}_a} \left(\Sigma_{\boldsymbol{S}_a,\mathcal{R}_a} \Sigma^T_{\boldsymbol{S}_a,\mathcal{R}_a} \left(\frac{\sigma^2_{\mathcal{R}_b}}{\sigma^4_{\mathcal{R}_a}} - \frac{1}{\sigma^2_{\mathcal{R}_a}} \right) + \Sigma_{\boldsymbol{S}_a} \right)^{-1} \Bigg)^T \Sigma_{\boldsymbol{S}_a,\mathcal{R}_a} \\
&= \frac{\sigma^4_{\mathcal{R}_b}}{\sigma^4_{\mathcal{R}_a}} \Sigma^T_{\boldsymbol{S}_a,\mathcal{R}_a} \left(\Sigma^{-1}_{S_a} - \frac{\left(\frac{\sigma^2_{\mathcal{R}_b}}{\sigma^4_{\mathcal{R}_a}} - \frac{1}{\sigma^2_{\mathcal{R}_a}}\right) \Sigma^{-1}_{S_a} \Sigma_{S_a,R_a} \Sigma^T_{S_a,R_a} \Sigma^{-1}_{S_a}}{1 + \left(\frac{\sigma^2_{\mathcal{R}_b}}{\sigma^4_{\mathcal{R}_a}} - \frac{1}{\sigma^2_{\mathcal{R}_a}}\right) \Sigma^T_{S_a,R_a} \Sigma^{-1}_{S_a} \Sigma_{S_a,R_a}} \right)^T \Sigma_{\boldsymbol{S}_a,\mathcal{R}_a} \\
&= \frac{\sigma^4_{\mathcal{R}_b}}{\sigma^4_{\mathcal{R}_a}} \left(OV_a - \frac{\left(\frac{\sigma^2_{\mathcal{R}_b}}{\sigma^4_{\mathcal{R}_a}} - \frac{1}{\sigma^2_{\mathcal{R}_a}}\right) OV_a^2}{1 + \left(\frac{\sigma^2_{\mathcal{R}_b}}{\sigma^4_{\mathcal{R}_a}} - \frac{1}{\sigma^2_{\mathcal{R}_a}}\right) OV_a} \right) \\
&= \frac{\sigma^4_{\mathcal{R}_b}}{\sigma^2_{\mathcal{R}_b} + \frac{\sigma^2_{\mathcal{R}_a}(\sigma^2_{\mathcal{R}_a} - OV_a)}{OV_a}}
\end{aligned}$$

where in the third row we used the Sherman-Morrison formula and the definition of OV_b. □

Remark. Notice that in the normal case, where both the returns distribution and signals are normally distributed, there is no loss of generality in assuming that high school graduates receive only a scalar noise of the form

$$s = \mathcal{R} + \epsilon$$

where $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$. Following the same steps as before, we can show that the observed variance of beliefs is given by

$$OV = \frac{\sigma_{\mathcal{R}}^4}{\sigma_{\mathcal{R}}^2 + \sigma_\epsilon^2}$$

which implies that the information structure $P(S|\mathcal{R}) = \mathcal{N}(\mathcal{R}, \sigma_\epsilon^2)$ is identified by

$$\sigma_\epsilon^2 = \frac{\sigma_{\mathcal{R}}^2(\sigma_{\mathcal{R}}^2 - OV)}{OV}$$

Given the information structure, the counterfactual distribution is simply given by

$$\frac{\sigma_a^4}{\sigma_a^2 - \frac{\sigma_{\mathcal{R},b}^2(\sigma_{\mathcal{R},b}^2 + OV_b)}{OV_b}}$$

which aligns with the counterfactual quantity when agents have a richer signal structure.

1.12.2.3 Identification of the Information Structure Decomposition with Data on the Full Belief Distribution

In some cases, researchers may hope to elicit information on the probabilities that an agent put on each outcome realization (Manski (2004), Wiswall and Zafar (2015), Zafar (2011), Wiswall and Zafar (2021), Diaz-Serrano and Nilsson (2022)). We now turn to show that this information is sufficient for point identification of our choice gap decomposition, with respect to the information structure.

We assume that individuals from group b have earnings distribution π_b and access to the information structure $(\mathrm{P}(S|b, \mathcal{R}), \mathcal{S})$, and for group a have returns distribution π_a and access to the information structure $(\mathrm{P}(S|a, \mathcal{R}), \mathcal{S})$. Denote by $q_{s,g} \in \Delta(\mathcal{R})$ the posterior beliefs induced by a signal realization $s \in S$ and prior π_g. We let $q_{s,g}(\mathcal{R})$ be the assigned density that this posterior puts on state $\mathcal{R}$. Furthermore, we assume that we observe for each group the joint distribution, $\phi(\mathcal{R}, q_s)$, of returns $\mathcal{R}$ and the posterior beliefs q_s.

We start by noting that within the framework, knowing beliefs allows us to identify a richer notion of costs. Specifically, denote by $B_i = \int_{\mathcal{R}} \mathcal{R} q_i(\mathbb{R}) d\mathcal{R}$ the measured posterior mean for individuals with beliefs q_i and notice that

$$P(D = 1|x, B) = E[\mathbb{1}[B_i \geq c(x, \nu)]] \tag{1.25}$$

where ν is additional cost heterogeneity, not included in our identifying discussion in section 1.2.5. Under some regularity conditions and the assumption $B \perp\!\!\!\perp \nu|X$, we can identify the distribution of $c(x, \nu)$ for each x and B, using variation in B. The identification here relies on B as a "special regressor" needed for identification, as discussed in (Lewbel (2012)). From now on we assume we know the joint distribution of $P(q_i, c_i|x)$, and omit the cost c.

To identify the outcomes distribution, we can use two approaches. The first is simply be able to observe the realization distribution if possible. The other is to use the measured beliefs and simply integrate over beliefs, i.e.

$$\pi_g(\mathcal{R}) = \int_i q_i(\mathcal{R}) di \tag{1.26}$$

Under the assumption that rational expectations are held, this should provide the initial prior.

We start by showing the following lemma that shows that for a fixed information structure, there's a mapping from the posterior, given prior π_g' to a posterior under a different prior. Let π_g and $\pi_{g'}$ be two priors with the same support, then for each s, information structure $P(s|\alpha)$ prior π_g and implied posterior q_s, the counterfactual posterior with prior $\pi_{g'}$ is given by $q_{s,g'} = \frac{\frac{q_s(\mathcal{R})}{\pi(\mathcal{R})} \pi_{g'}(\mathcal{R})}{\int_{\mathcal{R}} \frac{q_s(\mathcal{R})}{\pi(\mathcal{R})} \pi_{g'}(\mathcal{R}) d\mathcal{R}}$

Proof.

$$\begin{aligned} q_{s,g'}(\mathcal{R}) &= \frac{P(s|\mathcal{R})\pi_{g'}(\mathcal{R})}{\int_{\mathcal{R}} p(s|\mathcal{R})\pi_{g'}(\mathcal{R})d\mathcal{R}} \\ &= \frac{\frac{P(s)q_s(\mathcal{R})}{\pi(\mathcal{R})}\pi_{g'}(\mathcal{R})}{\int_{\mathcal{R}} \frac{P(s)q_s(\mathcal{R})}{\pi(\mathcal{R})}\pi_{g'}(\mathcal{R})d\mathcal{R}} \\ &= \frac{\frac{q_s(\mathcal{R})}{\pi(\mathcal{R})}\pi_{g'}(\mathcal{R})}{\int_{\mathcal{R}} \frac{q_s(\mathcal{R})}{\pi(\mathcal{R})}\pi_{g'}(\mathcal{R})d\mathcal{R}} \end{aligned}$$

□

Lemma 1.12.2.3 demonstrates that the counterfactual posterior can be calculated from the known posterior π_g and the counterfactual distribution $\pi_{g'}$, without requiring explicit knowledge of the information structure. Given the counterfactual posteriors, one can also derive the counterfactual means and thus identify all components of the decomposition. We proceed to establish that all parts of the decomposition are identified.

Recall that for our decomposition we needed to identify the distribution of counterfactual posterior mean, if the returns were drawn according to group b, information according to group a and updated correctly in this new counterfactual world.

$$P_{\langle a,b\rangle} = \int_{\mathcal{R}} \mathcal{P}(E_{a,b}(s) \geq 0|\mathcal{R}, a), \pi_b(\mathcal{R})d\mathcal{R}$$

Proposition 5. Assume we know $\phi_a(q_{s,a}, \boldsymbol{\mathcal{R}})$ and $\phi_b(q_{s,a}, \boldsymbol{\mathcal{R}})$ then the conditional distribution $\mathcal{P}(E_{a,b}(s)|\mathcal{R}, a)$ is identified and so is $P_{\langle a,b\rangle}$ in 1.22

Proof. The proof follows from Lemma 1.12.2.3. Notice that according to Lemma 1.12.2.3, every two signals that generate the same posterior for group a, also generate the same posteriors in the counterfactual case where $\mathcal{R}$ is distributed according to π_b; therefore, it's enough to know the posterior without requiring the information structure. Further, using Lemma 3, we can identify the distribution of the counterfactual posteriors by calculating the implied distribution of the composition $\frac{\left(\frac{q_s(\mathcal{R})}{\pi(\mathcal{R})}\right)\pi_{g'}(\mathcal{R})}{\int_{\mathcal{R}}\left(\frac{q_s(\mathcal{R})}{\pi(\mathcal{R})}\right)\pi_{g'}(\mathcal{R})d\mathcal{R}}$. Finally, to obtain $P(E_{a,b}|\mathcal{R})$, we only need to map each posterior to its implied mean. As $P(E_{a,b}|\mathcal{R})$ is identified, $P_{\langle a,b\rangle}$ is trivially identified, along with the decomposition components values. □

One implication of Proposition 5 is that in the case where we have binary outcomes $Y \in \{1, 0\}$, and we know the joint distribution of $\phi(E[Y|s], Y)$, the decomposition is point identified using simply the conditional mean beliefs. If outcomes are binary $Y \in \{1, 0\}$ and we observe the joint $\phi(E[Y|s], Y)$, then $P_{\langle a,b\rangle}$ in 1.22 is point identified

Proof. Simply follows from proposition 5 and the fact that in the bianry case $E[Y|s]$ is the posterior distribution. □

The case of binary outcomes is prevalent in many applications within the discrimination literature. For instance, in bail decisions, judges are often modeled as agents attempting to predict the likelihood of reoffense (e.g., Arnold et al. (2018)), Researchers may wish to quantify the extent to which disparities in decisions made for Black and White defendants are driven by the information available to judges or by the underlying distribution of reoffending rates. The above corollary demonstrates that we can decompose this gap and precisely identify the role each component plays. Similar arguments can be extended to other contexts, such as hiring decisions (Bertrand and Mullainathan (2004),Kline et al. (2022)) or treatment allocation in medical settings (Chan et al. (2022)).

1.12.3 Nonparametric Partial Identification

Researchers often have data on outcomes and posterior mean beliefs, accessible via the identification strategy outlined in Section 1.10.2 or through surveys querying individual beliefs. However, access to this data alone in general does not suffice for the point identification of the counterfactual beliefs distribution. Building on insights from the empirical information robustness literature(Bergemann and Morris (2019, 2013, 2016); Bergemann et al. (2022) Syrgkanis et al. (2017) Gualdani and Sinha (2019) Magnolfi and Roncoroni (2023)) we demonstrate a methodology to identify the counterfactual distribution of beliefs. Our proof in this section relies on Bergemann et al. (2022).

Our objective is to describe the identified set of the first and second parameter of interest. Following the last section, we assume that everything is conditioned on x, z, and subsume x and z for brevity, and assume to know the joint distribution $\phi(\mathcal{R}, \mathrm{E})$, for both grouaps a and b. Throughout the discussion we introduce and omit group membership when it's needed. Before we start, we redefine and define some of the notation we would be using in the discussion. We assume that individuals have access to information structure $\mathcal{S}$, with support s and density function $f(s|\mathcal{R})$ and the corresponding CDF $F(s|\mathcal{R})$. We denote by $\mu \in \Delta(\text{supp}(\mathcal{R}))$ the prior distribution. The posterior mean beliefs, given information structure $\mathcal{S}$ and prior μ is given by $E[\mathcal{R}|s; \mathcal{S}, \mu]$. Throughtout most of the discussion we would fix $\mathcal{S}$, and indicate it only when it matters. We further define the conditional distribution of beliefs, that are generated for a given prior and information structure, conditioned on $\mathcal{R}$ as

$$P^{\mu}_{\mathcal{S}}(\mathrm{E}|\mathcal{R}) = \int_{s:E[\mathcal{R}|s;\mathcal{S},\mu]=\mathrm{E}} dF(s|\mathcal{R})$$

Before moving to the main identification argument we show the following two trivial claims.

Claim 1. Let $\mathrm{E}(s,\mu)$ be

$$\mathrm{E}(s,\mu) = \arg\min_{\mathrm{E}} \int_{\mathcal{R}} (\mathcal{R}-\mathrm{E})^2 \frac{\mu(\mathcal{R})f(s|\mathcal{R})}{\int_{\mathcal{R}} \mu(\mathcal{R})f(s|\mathcal{R})d\mathcal{R}} d\mathcal{R} \tag{1.27}$$

then $\mathrm{E}(s,\mu) = E[\mathcal{R}|s;\mu]$

Proof. The results are simply implied by the first order conditions. □

Claim 2. Fix two prior distributions, $\mu, \mu' \in \Delta(\mathcal{R})$, where μ is absolute continuous with respect to μ' and let $\mathcal{E}(s)$ and $\mathcal{E}'(s)$ be

$$\mathcal{E}(s) = \arg\min_{\mathrm{E}} \int_{\mathcal{R}} (\mathcal{R}-\mathrm{E})^2 \frac{\mu(\mathcal{R})f(s|\mathcal{R})}{\int_{\mathcal{R}} \mu(\mathcal{R})f(s|\mathcal{R})d\mathcal{R}} d\mathcal{R} \tag{1.28}$$

and

$$\mathcal{E}'(s) = \arg\min_{\mathrm{E}} \int_{\mathcal{R}} \left[(\mathcal{R}-\mathrm{E})^2 \frac{\mu(\mathcal{R})}{\mu'(\mathcal{R})}\right] \frac{\mu'(\mathcal{R})f(s|\mathcal{R})}{\int_{\mathcal{R}} \mu'(\mathcal{R})f(s|\mathcal{R})d\mathcal{R}} d\mathcal{R} \tag{1.29}$$

Let $\Gamma(\mathcal{R},\mathcal{E})$ be the joint distribution of $\mathcal{R}$ and $\mathcal{E}(s)$, where $\Gamma(\mathcal{R},\mathcal{E};\mu) = \mu(\mathcal{R}) \int_{\{s:\mathcal{E}(s)=\mathcal{E}\}} dF(s|\mathcal{R})$, and let $\tilde{\Gamma}(\mathcal{R},\mathcal{E})$ be the joint distribution of $\mathcal{R}$ and $\mathcal{E}'(s)$, where $\tilde{\Gamma}(\mathcal{R},\mathcal{E}) = \mu'(\mathcal{R}) \int_{\{s:\mathcal{E}'(s)=\mathcal{E}\}} dF(s|\mathcal{R})$, then

$$\Gamma(\mathcal{R},\mathcal{E}) = \frac{\mu(\mathcal{R})}{\mu'(\mathcal{R})}\tilde{\Gamma}(\mathcal{R},\mathcal{E}) \tag{1.30}$$

Furthermore,

$$P_{\mathcal{S}}^{\mu}(\mathcal{E}|\mathcal{R}) = \frac{\tilde{\Gamma}(\mathcal{R},\mathcal{E})}{\mu'(\mathcal{R})} \tag{1.31}$$

Proof. Notice that for each signal realization $s \in S$ we have

$$\mathcal{E}'(s) = \arg\min_{\mathrm{E}} \int_{\mathcal{R}} (\mathcal{R} - \mathrm{E})^2 \frac{\mu(\mathcal{R})}{\mu'(\mathcal{R})} \frac{\mu'(\mathcal{R}) d(s|\mathcal{R})}{\int_{\mathcal{R}} \mu'(\mathcal{R}) d(s|\mathcal{R}) d\mathcal{R}} d\mathcal{R} \tag{1.32}$$

$$= \arg\min_{\mathrm{E}} \int_{\mathcal{R}} (\mathcal{R} - \mathrm{E})^2 \mu(\mathcal{R}) f(s|\mathcal{R}) d\mathcal{R} \tag{1.33}$$

$$= \mathcal{E}(s) \tag{1.34}$$

Therefore, $\int_{\{s:\mathcal{E}'(s)=\mathrm{E}\}} dF(s|\mathcal{R}) = \int_{\{s:\mathcal{E}(s)=\mathrm{E}\}} dF(s|\mathcal{R})$, for all E, which implies

$$\Gamma(\mathcal{R}, \mathcal{E}) = \mu(\mathcal{R}) \int_{\{s:\mathcal{E}(s)=\mathcal{E}\}} dF(s|\mathcal{R}) \tag{1.35}$$

$$= \mu(\mathcal{R}) \int_{\{s:\mathcal{E}'(s)=\mathcal{E}\}} dF(s|\mathcal{R}) \tag{1.36}$$

$$= \frac{\mu(\mathcal{R})}{\mu'(\mathcal{R})} \tilde{\Gamma}(\mathcal{R}, \mathcal{E}) \tag{1.37}$$

Finally, notice that $\mathcal{E}'(s) = \mathcal{E}(s) = E[\mathcal{R}|s; \mathcal{S}, \mu]$, therefore we have that

$$P^{\mu}_{\mathcal{S}}(\mathrm{E}|\mathcal{R}) = \int_{s:E[\mathcal{R}|s;\mathcal{S},\mu]=\mathrm{E}} dF(s|\mathcal{R}) = \int_{s:\mathcal{E}(s)=\mathrm{E}} dF(s|\mathcal{R}) = \int_{s:\mathcal{E}'(s)=\mathrm{E}} dF(s|\mathcal{R}) = \frac{\tilde{\Gamma}(\mathcal{R}, \mathrm{E})}{\mu'(\mathcal{R})} \tag{1.38}$$

□

We can now proceed to the identification argument. We want to describe the identified set of the information channel. The first component, which is observed share, is clearly identified, we therefore need only to show that the counterfactual share is identified. Fix an observed joint distribution of beliefs and states, induced by an unknown information structure $\mathcal{S}$ and μ', $\phi(\mathcal{R}, E) = \mu'(\mathcal{R}) P^{\mu'}_{\mathcal{S}}(E|\mathcal{R})$. We want to characterize the set of possible joint distributions of beliefs and states for the counterfactual case where we change the state distribution to μ, but leave the information structure $\mathcal{S}$ unchanged. Throught the discussion we assume that both priors have common support and that $\forall \mathcal{R} \mu(\mathcal{R}) \gg 0 \iff \mu'(\mathcal{R}) \gg 0$, such that our counterfactual would be well defined.

Denote by $\mathcal{C}(\phi(\mathcal{R}, \mathrm{E}), \mu')$ the set of joint distributions, $\tilde{\phi}(\mathcal{R}, \mathrm{E})$, that can be induced by the information structure $\mathcal{S}$, which induces $\phi(\mathcal{R}, \mathrm{E})$, and the returns distribution μ. i.e.

$$\mathcal{C}(\phi(\mathcal{R}, \mathrm{E}), \mu) = \Big\{ \tilde{\phi}(\mathcal{R}, \mathrm{E}) \in \Delta(\mathrm{supp}(\mathcal{R}), cl(\mathrm{supp}(\mathcal{R}))) \Big| \\ \exists \mathcal{S} \text{ s.t } \mu'(\mathcal{R}) P_{\mathcal{S}}^{\mu'}(\mathrm{E}|\mathcal{R}) = \phi(\mathcal{R}, \mathrm{E}), \mu(\mathcal{R}) P_{\mathcal{S}}^{\mu}(\mathrm{E}|\mathcal{R}) = \tilde{\phi}(\mathcal{R}, \mathrm{E}) \Big\}$$

where $cl(\mathrm{supp}(\mathcal{R}))$ is the support of beliefs. Our objective is to find a tractable characterization of this set. Let $\pi(\mathcal{R}, \mathrm{E}_{\mu'}, \mathrm{E}_{\mu}) \in \Delta(\mathcal{R}, cl(\mathcal{R}), cl(\mathcal{R}))$ be a joint distribution that satisfies

$$\int_{\mathrm{E}_{\mu}} \pi(\mathcal{R}, \mathrm{E}, \mathrm{E}_{\mu}) d\mathrm{E}_{\mu} = \phi(\mathcal{R}, \mathrm{E}) \tag{1.39}$$

$$\forall \mathrm{E}_{\mu'}, \mathrm{E}_{\mu} \quad \mathrm{E}_{\mu'} = \arg\min_{\mathrm{E}} \int_{\mathcal{R}} (\mathcal{R} - \mathrm{E})^2 \pi(\mathcal{R}, \mathrm{E}_{\mu'}, \mathrm{E}_{\mu}) d\mathcal{R} \tag{1.40}$$

$$\forall \mathrm{E}_{\mu'}, \mathrm{E}_{\mu} \quad \mathrm{E}_{\mu} = \arg\min_{\mathrm{E}} \int_{\mathcal{R}} (\mathcal{R} - \mathrm{E})^2 \frac{\mu(\mathcal{R})}{\mu'(\mathcal{R})} \pi(\mathcal{R}, \mathrm{E}_{\mu'}, \mathrm{E}_{\mu}) d\mathcal{R} \tag{1.41}$$

and denote the set of implied joint distribution of $\mathcal{R}$ and E_{μ} as

$$\mathcal{M}(\phi(\mathcal{R}, \mathrm{E}), \mu) = \Big\{ \tilde{\phi}(\mathcal{R}, \mathrm{E}) \in \Delta(\mathrm{supp}(\mathcal{R}), cl(\mathrm{supp}(\mathcal{R}))) \Big| \\ \tilde{\phi}(\mathcal{R}, \mathrm{E}) = \int_{\mathrm{E}_{\mu'}} \frac{\mu(\mathcal{R})}{\mu'(\mathcal{R})} \pi(\mathcal{R}, \mathrm{E}_{\mu'}, \mathrm{E}_{\mu}) d\mathrm{E}_{\mu'}, \pi \text{ satisfies } (1.39), (1.40), (1.41) \Big\}$$

Claim 3. For any observed distribution $\phi(\mathcal{R}, \mathrm{E}) \in \Delta(\mathrm{supp}(\mathcal{R}), cl(\mathrm{supp}(\mathcal{R})))$ and $\mu \in \Delta(\mathrm{supp}(\mathcal{R}))$

that is absolute continuous with respect to μ', we have

$$\mathcal{C}(\phi(\mathcal{R},\mathrm{E}),\mu) = \mathcal{M}(\phi(\mathcal{R},\mathrm{E}),\mu) \tag{1.42}$$

Proof. We start by showing that $\mathcal{M}(\phi(\mathcal{R},\mathrm{E}),\mu) \subseteq \mathcal{C}(\phi(\mathcal{R},\mathrm{E}),\mu)$. Let $\tilde{\phi}(\mathcal{R},\mathrm{E}) \in \mathcal{M}(\phi(\mathcal{R},\mathrm{E}),\mu)$ and let $\pi(\mathcal{R},\mathrm{E}_{\mu'},\mathrm{E}_{\mu})$ be the corresponding joint distribution that satisfies (1.39),(1.40),(1.41). Then define the the information structure $\mathcal{S}_{\mathrm{E}_{\mu'},\mathrm{E}_{\mu}}$ as

$$P(\mathrm{E}_{\mu'},\mathrm{E}_{\mu}|\mathcal{R}) = \frac{\pi(\mathcal{R},\mathrm{E}_{\mu'},\mathrm{E}_{\mu})}{\int \pi(\mathcal{R},\mathrm{E}_{\mu'},\mathrm{E}_{\mu})d(\mathrm{E}_{\mu'},\mathrm{E}_{\mu'})} = \frac{\pi(\mathcal{R},\mathrm{E}_{\mu'},\mathrm{E}_{\mu})}{\mu'(\mathcal{R})} \tag{1.43}$$

where the denominator follows from condition (1.39). Notice that as π satisfies condition (1.40), claim 1 and claim 2 implies

$$P^{\mu'}_{\mathcal{S}_{\mathrm{E}_{\mu'},\mathrm{E}_{\mu}}}(\mathrm{E}|\mathcal{R}) = \int_{\mathrm{E}_{\mu}} P(E,\mathrm{E}_{\mu}|\mathcal{R})d\mathrm{E}_{\mu} \tag{1.44}$$

hence, using constraint (1.39), we have

$$\mu'(\mathcal{R})P^{\mu'}_{\mathcal{S}_{\mathrm{E}_{\mu'},\mathrm{E}_{\mu}}}(\mathrm{E}|\mathcal{R}) = \mu'(\mathcal{R})\int_{\mathrm{E}_{\mu}} P(\mathrm{E},\mathrm{E}_{\mu}|\mathcal{R})d\mathrm{E}_{\mu} = \phi(\mathcal{R},\mathrm{E}) \tag{1.45}$$

Next, notice by constraint (1.41) and claim 2 we know that $\frac{\int_{\mathrm{E}_{\mu'}} \pi(\mathcal{R},\mathrm{E}_{\mu'},\mathrm{E}_{\mu})d\mathrm{E}_{\mu'}}{\mu'(\mathcal{R})} = P^{\mu}_{\mathcal{S}_{\mathrm{E}_{\mu'},\mathrm{E}_{\mu}}}(\mathrm{E}|\mathcal{R})$, then

$$\tilde{\phi}(\mathcal{R},\mathrm{E}) = \int_{\mathrm{E}_{\mu'}} \frac{\mu(\mathcal{R})}{\mu'(\mathcal{R})}\pi(\mathcal{R},\mathrm{E}_{\mu'},\mathrm{E}_{\mu})d\mathrm{E}_{\mu'} = \mu(\mathcal{R})P^{\mu}_{\mathcal{S}_{\mathrm{E}_{\mu'},\mathrm{E}_{\mu}}}(\mathrm{E}|\mathcal{R}) \tag{1.46}$$

therefore, we showed that there exist an information structure as needed, which implies $\tilde{\phi}(\mathcal{R},\mathrm{E}) \in$

$\mathcal{C}(\phi(\mathcal{R},\mathrm{E}),\mu)$

To see the reverse inclusion, $\mathcal{C}(\phi(\mathcal{R},\mathrm{E}),\mu) \subseteq \mathcal{M}(\phi(\mathcal{R},\mathrm{E}),\mu)$. Fix $\tilde{\phi}(\mathcal{R},\mathrm{E}) \in \mathcal{C}(\phi(\mathcal{R},\mathrm{E}),\mu)$ and let $\mathcal{S}$ be the information structure that satisfies

$$\mu'(\mathcal{R}) P^{\mu'}_{\mathcal{S}}(\mathrm{E}|\mathcal{R}) = \phi(\mathcal{R},\mathrm{E}) \tag{1.47}$$

$$\mu(\mathcal{R}) P^{\mu}_{\mathcal{S}}(\mathrm{E}|\mathcal{R}) = \tilde{\phi}(\mathcal{R},\mathrm{E}) \tag{1.48}$$

Define the functions $\mathrm{E}_\mu : S \to cl(Supp(\mathcal{R})), \mathrm{E}'_\mu : S \to cl(Supp(\mathcal{R}))$ as

$$\mathrm{E}_{\mu'}(s) = \arg\min_{\mathrm{E}} \int_{\mathcal{R}} (\mathcal{R}-\mathrm{E})^2 \mu'(\mathcal{R}) f(s|\mathcal{R}) d\mathcal{R} \tag{1.49}$$

$$\mathrm{E}_{\mu}(s) = \arg\min_{\mathrm{E}} \int_{\mathcal{R}} (\mathcal{R}-\mathrm{E})^2 \frac{\mu(\mathcal{R})}{\mu'(\mathcal{R})} \mu'(\mathcal{R}) f(s|\mathcal{R}) d\mathcal{R} \tag{1.50}$$

and define the joint probability $\pi(\mathcal{R},\mathrm{E}_{\mu'},\mathrm{E}_\mu)$ as

$$\pi(\mathcal{R},\mathrm{E}_{\mu'},\mathrm{E}_\mu) = \mu'(\mathcal{R}) \int_{s:\mathrm{E}_{\mu'}(s)=\mathrm{E}_{\mu'},\mathrm{E}_\mu(s)=\mathrm{E}_\mu} dF(s|\mathcal{R}) \tag{1.51}$$

Next, using claim 2 we know that $\mathrm{E}'_\mu(s) = E[\mathcal{R}|s;\mathcal{S},\mu']$ and therefore

$$\int_{\mathrm{E}_\mu} \pi(\mathcal{R},\mathrm{E},\mathrm{E}_\mu) d\mathrm{E}_\mu = \mu'(\mathcal{R}) \int_{\mathrm{E}_\mu} \pi(\mathrm{E},\mathrm{E}_\mu|\mathcal{R}) d\mathrm{E}_\mu = \mu'(\mathcal{R})\pi(\mathrm{E}|\mathcal{R}) = \mu'(\mathcal{R}) P^{\mu'}_{\mathcal{S}}(\mathrm{E}|\mathcal{R}) = \phi(\mathcal{R},\mathrm{E}) \tag{1.52}$$

To see that π satisfies condition (1.40), we can use the law of iterated expectations

$$
\begin{aligned}
\forall \mathrm{E}_{\mu'}, \mathrm{E}_{\mu} \quad & \arg\min_{\mathrm{E}} \int_{\mathcal{R}} (\mathcal{R} - \mathrm{E})^2 \pi(\mathcal{R}, \mathrm{E}_{\mu'}, \mathrm{E}_{\mu}) d\mathcal{R} && (1.53)\\
&= \arg\min_{\mathrm{E}} \int_{\mathcal{R}} (\mathcal{R} - \mathrm{E})^2 \int_{s} \pi(\mathcal{R}, \mathrm{E}_{\mu'}, \mathrm{E}_{\mu}, s) ds d\mathcal{R} && (1.54)\\
&= \arg\min_{\mathrm{E}} \int_{s:\mathrm{E}_{\mu'}(s)=\mathrm{E}_{\mu'}, \mathrm{E}_{\mu}(s)=\mathrm{E}_{\mu}} \int_{\mathcal{R}} (\mathcal{R} - \mathrm{E})^2 \pi(\mathcal{R}, \mathrm{E}_{\mu'}, \mathrm{E}_{\mu}, s) && (1.55)\\
&= \mathrm{E}_{\mu'} && (1.56)
\end{aligned}
$$

where we used the fact that E_{μ}' minimizes the expression by construction. A similar argument shows that (1.41) also holds. Finally, by claim 2, condition (1.41), and the way π is constructed, we have that

$$
\int_{\mathrm{E}_{\mu'}} \frac{\mu(\mathcal{R})}{\mu'(\mathcal{R})} \pi(\mathcal{R}, \mathrm{E}_{\mu'}, \mathrm{E}_{\mu}) d\mathrm{E}_{\mu'} = \mu(\mathcal{R}) P_{\mathcal{S}}^{\mu}(\mathrm{E}|\mathcal{R}) = \tilde{\phi}(\mathcal{R}, \mathrm{E}) \tag{1.57}
$$

which implies that $\tilde{\phi}(\mathcal{R}, \mathrm{E}) \in \mathcal{M}(\phi(\mathcal{R}, \mathrm{E}), \mu)$ □

To conclude the identification argument, we introduce the following assumption:

Assumption 5. μ_a *is absolutely continuous with respect to* μ_b.

We fix cost c, and denote the set of the possible probabilities

$$
P_{\langle a,b \rangle}(c) = \mathcal{P}(E_{a,b} \geq c | \mathcal{R}, c, a) \pi_b(\mathcal{R}|c) \tag{1.58}
$$

as

$$\mathcal{I}(\phi_a(\mathcal{R},\mathrm{E})) = \left\{ p \in [0,1] \middle| p = \int_{\mathcal{R}} \int_{\mathrm{E}\geq c} \tilde{\phi}(\mathcal{R},\mathrm{E}) d\mathrm{E} d\mathcal{R} \right.$$
$$\left. \text{s.t } \tilde{\phi}(\mathcal{R},\mathrm{E}) \in \mathcal{C}(\phi_a(\mathcal{R},\mathrm{E}), \mu_b(\mathcal{R})) \right\}$$

The following claim shows an easy characterization of this set

Claim 4. The identified set is given by

$$\mathcal{I}(\phi_a(\mathcal{R},\mathrm{E})) = \left\{ p \in [0,1] \middle| p = \int_{\mathcal{R}} \int_{\mathrm{E}\geq c} \tilde{\phi}(\mathcal{R},\mathrm{E}) d\mathrm{E} d\mathcal{R} \right.$$
$$\left. \text{s.t } = \tilde{\phi}(\mathcal{R},\mathrm{E}) \in \mathcal{M}(\phi(\mathcal{R},\mathrm{E}_{\eta_a}), \mu_b(\mathcal{R})) \right\}$$

Proof. Follows from claim 3 and assumption 5. □

Proposition 6. The quantity in 1.22 is partially identified given the distribution of $\phi(\mathcal{R},c,\mathrm{E})$

Proof. follows trivially from claim 4. □

Notice that we can further simplify the characterization of the identified set by using the fact that constraint (1.40) and (1.41) are satisfied if and only if the first order conditions hold. Therefore, we can rewrite the constraints (1.40) and (1.41) as

$$\forall \mathrm{E}_{\mu'}, \mathrm{E}_{\mu} \quad \mathrm{E}_{\mu'} = \int_{\mathcal{R}} (\mathcal{R} - \mathrm{E}) \pi(\mathcal{R}, \mathrm{E}_{\mu'}, \mathrm{E}_{\mu}) d\mathcal{R} \tag{1.40a}$$

$$\forall \mathrm{E}_{\mu'}, \mathrm{E}_{\mu} \quad \mathrm{E}_{\mu} = \int_{\mathcal{R}} (\mathcal{R} - \mathrm{E}) \frac{\mu(\mathcal{R})}{\mu'(\mathcal{R})} \pi(\mathcal{R}, \mathrm{E}_{\mu'}, \mathrm{E}_{\mu}) d\mathcal{R} \tag{1.41a}$$

Now, as the constraints (1.39), (1.40a) and (1.41a) are linear, the identified set is convex, and we can define it as an interval bounded between $[\underline{p}, \overline{p}]$, such that

$$\underline{p}, \overline{p} = \min_{\pi}, \max_{\pi} \int_z h_a(z) \int_{\mathcal{R}} \int_{\mathrm{E}_a} \int_{\mathrm{E}_b \geq c(z)} \pi(\mathcal{R}, \mathrm{E}_a, \mathrm{E}_b) \frac{\mu_b(\mathcal{R})}{\mu_a(\mathcal{R})} d(\mathcal{R}, \mathrm{E}_a, \mathrm{E}_b, z) \tag{1.59}$$

s.t

$$\forall \mathrm{E}, \mathcal{R} \quad \phi(\mathcal{R}, \mathrm{E}) = \int \pi(\mathcal{R}, \mathrm{E}, \mathrm{E}_b) d\mathrm{E}_b \tag{1.60}$$

$$\forall \mathrm{E}_a, \mathrm{E}_b \quad \mathrm{E}_a = \int (\mathcal{R} - \mathrm{E}_a) \pi(\mathcal{R}, \mathrm{E}_a, \mathrm{E}_b) d\mathcal{R} \tag{1.61}$$

$$\forall \mathrm{E}_a, \mathrm{E}_b \quad \mathrm{E}_b = \int (\mathcal{R} - \mathrm{E}_b) \frac{\mu_b(\mathcal{R})}{\mu_a(\mathcal{R})} \pi(\mathcal{R}, \mathrm{E}_a, \mathrm{E}_b) d\mathcal{R} \tag{1.62}$$

1.13 The Texas Higher Education Opportunity Project (THEOP)

The Texas Higher Education Opportunity Project (THEOP) is a comprehensive study designed to evaluate college planning and enrollment patterns in the context of Texas's policy granting automatic admission to public colleges and universities for students graduating in the top decile of their high school class. Collecting data from nine diverse Texas colleges and universities, including both public and private institutions, THEOP encompasses administrative records on applications, admissions, and enrollments, alongside a longitudinal survey of students from two cohorts in 2002. The administrative data set includes College Application Data, tracking demographics, academic profiles, and admission outcomes from before and after the 1998 implementation of the top 10% law, and College Transcript Data, detailing academic performance and progress of enrolled students. Efforts to ensure data quality and confidentiality have been meticulously undertaken, involving the removal of personal identifiers and the adjustment of data to prevent identification, thus ensuring a high level of privacy and data integrity. For our purpose we use questions asked about the interaction between high school students and their school councilor.

References

Abdulkadiroğlu, A., Pathak, P. A., Schellenberg, J., and Walters, C. R. (2020). Do parents value school effectiveness? *American Economic Review*, 110(5):1502–1539.

Adams-Prassl, A. and Andrew, A. (2019). Preferences and beliefs in the marriage market for young brides.

Agency, T. E. (2023). Enrollment in texas public schools 2022-23.

Angrist, J. D. and Imbens, G. W. (1995). Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475.

Arnold, D., Dobbie, W., and Hull, P. (2022). Measuring racial discrimination in bail decisions. *American Economic Review*, 112(9):2992–3038.

Arnold, D., Dobbie, W., and Yang, C. S. (2018). Racial bias in bail decisions. *The Quarterly Journal of Economics*, 133:1885.

Ben-Porath, Y. (1967). The production of human capital and the life cycle of earnings. *Journal of Political Economy*, 75(4):352–365.

Bergemann, D., Brooks, B., and Morris, S. (2022). Counterfactuals with latent information. *American Economic Review*, 112(1):343–368.

Bergemann, D. and Morris, S. (2013). Robust predictions in games with incomplete information. *Econometrica*, 81(4):1251–1308.

Bergemann, D. and Morris, S. (2016). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522.

Bergemann, D. and Morris, S. (2019). Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95.

Bertrand, M. and Mullainathan, S. (2004). Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American Economic Review*, 94(4):991–1013.

Blackwell, D. (1953). Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272.

Blinder, A. S. (1973). Wage discrimination: Reduced form and structural estimates. *The Journal of Human Resources*, 8(4):436–455.

Bohren, J. A., Hull, P., and Imas, A. (2022). Systemic discrimination: Theory and measurement. Technical Report w29820, National Bureau of Economic Research.

Borghans, L., Golsteyn, B. H., Heckman, J. J., and Humphries, J. E. (2016). What grades and achievement tests measure. *Proceedings of the National Academy of Sciences*, 113(47):13354–13359.

Brave, S. and Walstrum, T. (2014). Estimating marginal treatment effects using parametric and semiparametric methods. *The Stata Journal*, 14(1):191–217.

Cameron, S. V. and Taber, C. (2004). Estimation of educational borrowing constraints using returns to schooling. *Journal of Political Economy*, 112(1):132–182.

Canay, I. A., Mogstad, M., and Mountjoy, J. (2020). On the use of outcome tests for detecting bias in decision making. Technical Report w27802, National Bureau of Economic Research.

Caplin, A., Dean, M., and Leahy, J. (2022). Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy*, 130(6):1676–1715.

Card, D. (1995). Using geographic variation in college proximity to estimate the return to schooling. In Christofides, L., Grant, K., and Swidinsky, R., editors, *Aspects of Labour Market Behaviour: Essays in Honour of John Vanderkamp*, pages 201–222. University of Toronto Press, Toronto.

Carneiro, P., Heckman, J. J., and Vytlacil, E. J. (2011). Estimating marginal returns to education. *American Economic Review*, 101(6):2754–2781.

Carneiro, P. and Lee, S. S. (2009). Estimating distributions of potential outcomes using local instrumental variables with an application to changes in college enrollment and wage inequality. *Journal of Econometrics*, 149(2):191–208.

Chan, D. C., Gentzkow, M., and Yu, C. (2022). Selection with variation in diagnostic skill: Evidence from radiologists. *The Quarterly Journal of Economics*, 137(2):729–783.

Coate, S. and Loury, G. C. (1993). Will affirmative-action policies eliminate negative stereotypes? *The American Economic Review*, 83(5):1220–1240.

Cunha, F. and Heckman, J. (2007a). The technology of skill formation. *American Economic Review*, 97(2):31–47.

Cunha, F. and Heckman, J. J. (2007b). Identifying and estimating the distributions of ex post and ex ante returns to schooling. *Labour Economics*, 14(6):870–893.

Cunha, F., Nielsen, E., and Williams, B. (2021). The econometrics of early childhood human capital and investments. *Annual Review of Economics*, 13:487–513.

d'Haultfoeuille, X. and Maurel, A. (2013). Inference on an extended roy model, with an application to schooling decisions in france. *Journal of Econometrics*, 174(2):95–106.

Diaz-Serrano, L. and Nilsson, W. (2022). The reliability of students' earnings expectations. *Labour Economics*, 76:102182.

Eisenhauer, P., Heckman, J. J., and Vytlacil, E. (2015). The generalized roy model and the cost-benefit analysis of social programs. *Journal of Political Economy*, 123(2):413–443.

Evdokimov, K. and White, H. (2012). Some extensions of a lemma of kotlarski. *Econometric Theory*, 28(4):925–932.

Freeman, R. (1971). *The Market for College-Trained Manpower*. Harvard University Press, Cambridge.

Garip, F. (2020). What failure to predict life outcomes can teach us. *Proceedings of the National Academy of Sciences*, 117(15):8234–8235.

Gentzkow, M. and Kamenica, E. (2016). A rothschild-stiglitz approach to bayesian persuasion. *American Economic Review*, 106(5):597–601.

Gilraine, M., Gu, J., and McMillan, R. (2020). A new method for estimating teacher value-added. Technical Report w27094, National Bureau of Economic Research.

Giustinelli, P. (2022). Expectations in education: Framework, elicitation, and evidence. Available at SSRN 4318127.

Gualdani, C. and Sinha, S. (2019). Identification and inference in discrete choice models with imperfect information.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2nd edition.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica: Journal of the Econometric Society*, pages 153–161.

Heckman, J. J. and Robb, R. J. (1985). Alternative methods for evaluating the impact of interventions: An overview. *Journal of Econometrics*, 30(1-2):239–267.

Heckman, J. J. and Vytlacil, E. (2005). Structural equations, treatment effects, and econometric policy evaluation. *Econometrica*, 73(3):669–738.

Hull, P. (2021). What marginal outcome tests can tell us about racially biased decision-making. Working Paper w28503, National Bureau of Economic Research.

Kane, T. J. and Rouse, C. E. (1995). Labor-market returns to two- and four-year college. *American Economic Review*, 85(3):600–614.

Kapor, A. (2020). Distributional effects of race-blind affirmative action. Technical Report 2020-51, Unspecified Institution.

Kitagawa, E. M. (1955). Components of a difference between two rates. *Journal of the American Statistical Association*, 50(272):1168–1194.

Kline, P., Rose, E. K., and Walters, C. R. (2022). Systemic discrimination among large us employers. *The Quarterly Journal of Economics*, 137(4):1963–2036.

Kling, J. R. (2001). Interpreting instrumental variables estimates of the returns to schooling. *Journal of Business & Economic Statistics*, 19(3):358–364.

Lewbel, A. (2012). An overview of the special regressor method. Technical Report 810, Boston College Working Papers in Economics.

Lundberg, S. J. and Startz, R. (1983). Private discrimination and social intervention in competitive labor markets. *The American Economic Review*, 73(3):340–347.

Magnolfi, L. and Roncoroni, C. (2023). Estimation of discrete games with weak assumptions on information. *The Review of Economic Studies*, 90(4):2006–2041.

Manski, C. F. (2004). Measuring expectations. *Econometrica*, 72(5):1329–1376.

Matzkin, R. L. (1992). Nonparametric identification and estimation of polychotomous choice models. *Journal of Econometrics*, 58(1-2):137–168.

Matzkin, R. L. (1993). Nonparametric and distribution-free estimation of the binary threshold crossing and the binary choice models. *Econometrica*, pages 239–270.

Maćkowiak, B., Matějka, F., and Wiederholt, M. (2023). Rational inattention: A review. *Journal of Economic Literature*, 61(1):226–273.

Mountjoy, J. (2022). Community colleges and upward mobility. *American Economic Review*, 112(8):2580–2630.

Murnane, R. J., Willett, J. B., Duhaldeborde, Y., and Tyler, J. H. (2000). How important are the cognitive skills of teenagers in predicting subsequent earnings? *Journal of Policy Analysis and Management*, 19(4):547–568.

Neal, D. A. and Johnson, W. R. (1996). The role of premarket factors in black-white wage differences. *Journal of Political Economy*, 104(5):869–895.

Nybom, M. (2017). The distribution of lifetime earnings returns to college. *Journal of Labor Economics*, 35(4):903–952.

Oaxaca, R. (1973). Male-female wage differentials in urban labor markets. *International Economic Review*, 14(3):693–709.

Rao, C. R. (1992). *Linear Statistical Inference and its Applications*. Wiley, 2nd edition.

Roy, A. B. (1951). Some thoughts on the distribution of earnings. *Oxford Economic Papers*, 3(2):135–146.

Salganik, M. J. et al. (2020). Measuring the predictability of life outcomes with a scientific mass collaboration. *Proceedings of the National Academy of Sciences*, 117(15):8398–8403.

Small, M. L. and Pager, D. (2020). Sociological perspectives on racial discrimination. *Journal of Economic Perspectives*, 34(2):49–67.

Syrgkanis, V., Tamer, E., and Ziani, J. (2017). Inference on auctions with weak assumptions on information. *arXiv preprint arXiv:1710.03830*.

Times, T. N. Y. (2014). Where we came from, state by state. Accessed: 2023-11-01.

Vytlacil, E. (2002). Independence, monotonicity, and latent index models: An equivalence result. *Econometrica*, 70(1):331–341.

Vytlacil, E. (2006). A note on additive separability and latent index models of binary choice: Representation results. *Oxford Bulletin of Economics and Statistics*, 68(4):515–518.

Walters, C. R. (2018). The demand for effective charter schools. *Journal of Political Economy*, 126(6):2179–2223.

Watts, T. W. (2020). Academic achievement and economic attainment: Re-examining associations between test scores and long-run earnings. *AERA Open*.

Willis, R. J. and Rosen, S. (1979). Education and self-selection. *Journal of Political Economy*, 87(5, Part 2):S7–S36.

Wiswall, M. and Zafar, B. (2015). How do college students respond to public information about earnings? *Journal of Human Capital*, 9(2):117–169.

Wiswall, M. and Zafar, B. (2021). Human capital investments and expectations about career and family. *Journal of Political Economy*, 129(5):1361–1424.

Zafar, B. (2011). How do college students form expectations? *Journal of Labor Economics*, 29(2):301–348.

CHAPTER 2

IT'S NOT WHO YOU ARE, IT'S WHAT THEY KNOW: WAGE GAPS AND INFORMATIONAL FRICTIONS

2.1 Introduction

Firms can rarely, if ever, hire a worker after obtaining complete information on the worker's productivity and their outside options. These informational frictions can arise from various sources, including an inefficient hiring process, imperfect information on the firm's own production technology, attention costs of the interviewer, or cognitive biases on their part. These frictions have proven to be of great importance in countless theoretical results (e.g., Aigner and Cain (1977), Phelps (1972), Spence (1978), Bergemann et al. (2021b), Bergemann and Morris (2019a)). However, while we would like to take these into account both in modeling firms' decisions and in empirical exercises, this has proven to be fairly difficult. There are many informational environments in which firms operate that are not observable by researchers. Despite the crucial place information holds in theoretical research, most of the empirical literature on wage gaps has focused on other differences between groups. These differences are brought on by the structure of the labor market. The first type of fundamental differences is driven by the workers' productivity distribution. As many papers argue (for example, Altonji and Blank (1999), Blau and Kahn. (2017), Goldin (2014)), differences in workers' abilities between groups can drive differences in observed wages. These differences can stem from various sources, such as pre-market conditions, which generate differences in workers' productivity. Other sources that have been considered extensively include firms' taste-based discrimination, which can affect firms' willingness to pay for workers of different groups, or self-selection of workers into different occupations and industries due to differences in preferences, to name a few. The second type of mechanism that has been thoroughly explored in the literature and can be used as a cause for differences in wages is market frictions, such as search costs, probability of finding a job, differences in bargaining power, and differences in outside options.

In contrast to these explanations, this paper explores whether differences in wage distributions can stem from differences in firms' information about workers' productivity and their potential outside options. These differences can potentially be significant, affecting the wage distributions of different groups in various way, creating different in wage prospect for workers who are productively equivalent.

To gauge the potential importance of information frictions, we construct a static, parsimonious, common-value auction model of the labor market. This model assumes that heterogeneous workers receive job offers from identical firms that differ only in the information available to them. We then explore how this varies across markets with different levels of search frictions, captured by the number of wage offers a single worker receives. Since we are interested in examining the impact of information on the labor market, we leave unspecified the information firms have. We then ask how much of the wage gap between workers can be explained by the correlation between gender and race and the other information firms observe before making a wage offer. To form this test, we leverage an equivalence result from the robust prediction literature and information design (Bergemann and Morris (2013), Bergemann and Morris (2016)). This result shows that the set of distribution outcomes that can arise under a Bayes Nash Equilibrium (BNE) with some information structure corresponds to a set of joint distributions of actions and states, known as Bayes Correlated Equilibrium (BCE). We use this equivalence result to partially identify the set of possible productivity distributions which, with some signal structure possibly correlated with race and gender, can give rise to the observed wage gaps.

We find that information can potentially have a very large effect on the wage distribution, creating a significant divergence between workers' marginal product and their wage. For example, without any assumptions on the information available to firms, and in markets with relatively low search frictions, we can bound the mean productivity of white male workers to be between approximately $48,000 per year, which is roughly their mean wage, and $128,500. Moreover, we find that we can explain all of the wage gaps between white men, white women, and black men and women without needing to assume differences in productivity. More specifically, we find that the entire wage gap in our sample can be attributed to information frictions and can be supported by a productivity distribution with a mean bounded between approximately $48,000 and $132,800 per year in an economy with large search frictions, and between $48,000 and $93,600 in an economy with low search frictions.

This paper contributes to the vast literature on discrimination and, specifically, on statistical discrimination (Arrow (1973), Phelps (1972), Aigner and Cain (1977), Altonji and Pierret (2001), Lange (2007)). These early papers show that different information can give rise to differences in wage distributions but, as we argue in section 2.2, fail to explain differences in average wage. To address this issue, follow-up papers by Lundberg and Startz (1983) and Coate and Loury (1993) offer models in which minority workers end up investing less in human capital, generating equilibrium differences in workers' productivity available to firms. Unlike previous papers that attempt to explain wage gaps using statistical discrimination, we ask whether gaps can be explained without needing to change the underlying distribution of workers' productivity, but by relaxing the assumptions on the types of information firms have and allowing for firms to act based on private information. While our model does not exclude taste-based discrimination, it assumes that it's another force that affects the productivity distribution of workers as seen from the firms' perspective. A recent paper, by Chambers and Echenique (2021), also spotlight information frictions. They explore whether wage gaps and discriminatory policies can potentially arise from differences in information in an environment where the same worker can generate varying productivity for different firms possessing distinct information. Theoretically, they find that such an environment invariably supports an information structure capable of creating wage differences. In contrast to their work, our paper focuses on an empirical exercise involving homogeneous firms with different information on workers, all of whom can generate the same surplus for these firms. Our findings underscore that a key driver of divergent wage policies is the may be information itself, rather than the underlying productivity distribution, which remains constant in our exercise.

As discussed above, this paper builds on recent results from the robust prediction literature (Bergemann and Morris (2013), Bergemann and Morris (2016), Bergemann et al. (2017)) that explore the range of outcomes that can arise in a game with an unspecified information structure. These results are being used for informationally robust identification in a growing number of papers. Syrgkanis et al. (2021) is the most similar paper to ours; it explores how to achieve identification in a model of general second and first-price auctions without specifying the information available to individuals. They then use their identification results to analyze second-price auctions in the BingeAds sponsored auction marketplace and the OCS auction dataset to infer the underlying valuation distributions. Magnolfi and Roncoroni (2017) uses the BCE in an entry game with binary actions to identify the set of parameters on the utility function that are robust to the information firms have. Gualdani and Sinha (2019) employs the BCE framework we work with in this paper to identify the set of parameters and distributions governing an agent in a discrete choice model without specifying the information structure. Finally, Bergemann et al. (2021a) consider how to perform counterfactual analysis while holding the information fixed.

Additionally, this paper contributes to the recent empirical literature that emphasizes the role of workers' outside options in wage gaps. Caldwell and Danieli (2021) uses a two-sided matching model with transfers, based on Shapley and Shubik (1971), where heterogeneous workers and firms have idiosyncratic preferences for each other. They then calculate their outside option index for workers in Germany and find that it can explain roughly 25% of the wage gap. Black (1995) constructs a search model where some discriminatory employers reduce workers' outside options, thereby generating a wage decrease. In our model, the distribution of outside options is not generated by assuming workers have different preferences or due to monopolistic power, but because of the information firms have on workers and other firms' willingness to pay. Compared to previous papers, we allow for uncertainty over workers' outside options.

The paper proceeds as follows: Section 2 introduces the model. Section 3 discusses identification and how to test for the potential role of information in shaping the wage gap. Section 4 focuses on inference and computation, Section 5 presents our data and results, and Section 6 concludes.

2.2 A Simple Common-Value Auction Model of the Labour Market

2.2.1 The model

Let $\mathcal{J}$ be the set of firms in the market. Let $\mathcal{I}$ be the set of workers. There are $|\mathcal{G}|$ groups of workers. Workers have heterogeneous productivity, $v \in \mathcal{V} \subset \mathbb{R}_+$, drawn from a distribution $\mu(v|g) \in \Delta(\mathcal{V})$. A job offer from firm j to worker i consists of a wage $w_j^i \in W$. We assume that workers receive $N \leq |\mathcal{J}|$ jobs offers. We denote as $\boldsymbol{w}_i \in \mathcal{W} = W^N$ the vector of wage offers worker i receives. We further assume that both firms and workers are risk neutral and that the firms' production function is additive in the number workers. Therefore, if a firm succeeds in hiring a worker, that firm's marginal profit is given by $v - w$. We assume firms don't have a cost of making a wage offer. Worker i's utility from a vector of wage offers $\boldsymbol{w}_i$ is $u(\boldsymbol{w}) = \max_j \boldsymbol{w_i}$. Implying workers choose to worker at the firm who offers the highest wage. In the case of a tie, the worker selects at randiom one of the firms who offers the highest wage.

Before extending a wage offer, we assume that all firms observe both the worker's group, $g_i \in \mathcal{G}$, and a public signal, $x_i \in \mathcal{X}$, observed by all firms and by the econometrician. We do not restrict the correlation between the public signal and the workers' productivity. In addition to these signal, firms may observe additional signal, possibly private, $t_j \in \mathcal{T}_j$, prior to making a wage offer to the worker. The signal vector $t = (t_1, ..., t_J)$ can be arbitrarily correlated with both the worker's productivity and the public signals. We also do not put any restriction on the correlations between the different firms' signals. We denote the augmented signal structure $(\mathcal{G}, \mathcal{X}, \mathcal{T}, \mathrm{P}(g, x, t|v))$ by $S \in \mathcal{S}$. Let $k_i(\boldsymbol{w}) = \arg\max_j \boldsymbol{w}$ be the set of firms that offer the highest wage to worker i, then the worker is allocated to firm j with probability

$$q_j^i(\boldsymbol{w}) = \begin{cases} \frac{1}{|k(\boldsymbol{w})|} & \text{if j} \in k(\boldsymbol{w}) \\ 0 & otherwise \end{cases}$$

Finally, firms' j interim-expected marginal profit from offering a wage w_j to worker i, after observing the worker's public signals x_i, g_i and the private signal t_j is

$$\mathrm{E}\Big[(v_i - w_j)q(\boldsymbol{w})|t_j, x_i, g_i\Big] \propto$$
$$\sum_v \sum_{t_{-j}} \sum_{\boldsymbol{w}_{-j}} (v_i - w_j) q_j^i(\boldsymbol{w}) \Big[\prod_{k \neq j} \beta_k(w_k|t_k, x_i, g_i)\Big] p(t|v_i, x_i, g_i)\mu(v_i|x_i, g_i)p(g_i, x_i)$$

where $\beta_k(w|.)$ is the wage policy functions of firm k, given the firm's signals. A Bayes Nash Equilibrium (BNE) in this model is a mapping $\beta_k : \mathcal{S} \to \Delta(W_k)$ for each firm j, such that the firm maximizes its expected profit, conditional on their signals, and that workers choose to work at the firm that offered the highest wage.

In the model above, heterogeneity in wage offers stems from firms having access to different information. Specifically, information in the model plays two key roles in determining wage. The first is by affecting the firm's evaluation of worker productivity. Firms having different information structures implies that different firms evaluate the worker productivity differently, affecting their willingness to pay. The second channel through which information affects wages is firms' belief on the worker's outside option. Within the model, firms are asking themselves what other firms know about the worker and try to guess what other firms would be willing to pay for the worker.

To see the importance of these two channels, consider a simple setup with two firms, in which worker productivity is distributed uniformly between 0 and 1. Assume that that the two firms' signals are perfectly correlated. In that case, each firm knows that the other firm observes the same signal, then they would end up conducting a Bertrand competition, where wages would be the expected worker's productivity, given the common signal the average wage would be the worker's average productivity, as discussed in section 2.2.2. On the other hand, consider the polar opposite case, in which we have two firms, one is uninformed while the other one is perfectly informed.[1] An equilibrium in this setup would be that the informed firm would offer a wage of $\frac{v}{2}$, while the uninformed firm would randomize between 0 and 0.5 and the average wage would be $\frac{1}{3}$. To see this, notice that the uninformed firm would never make an offer higher than $\frac{1}{2}$, as it has negative ex-ante surplus. Next, notice that for any wage offer $w_{UI} \in [0, 0.5]$ the uninformed firm makes, the workers productivity, conditioned on the uniformed making the higher offer, is distributed uniformly between $[0, 2w_{UI}]$, and therefore, the expected surplus of the uninformed firm is $E[v - w_{UI}|\text{UI wins}] = 0$ for any $w_{UI} \in [0, 0.5]$. Finally,

1. This example is taken from Milgrom and Weber (1982)

the informed firm surplus from offering wage w_I is given by $(v - w_I)P(w_I > w_{UI}) = (v - w_I)\frac{w_I}{0.5}$. which is maximized at $w_I = \frac{v}{2}$. Given these two equilibrium strategies, the worker's average wage would be $\frac{1}{3}$,[2] which is lower than the average wage under first case, or under complete information. Therefore, we can see that simply changing the firms' access to information may have a large effect on the realized wage distribution and on the relation between workers productivity and their wages.

2.2.2 Relation To Phelps (1972)

Before moving forward and considering a whether a general information structure might be needed to explain wage gaps, we can first ask whether there exists a public signal, available to all firms, which can induce the observed wage distributions of workers from two groups. In his seminal paper on statistical discrimination, Phelps (1972) considers a model similar to ours, but restricts attention to public and normal signals. In his model, there exist incomplete information on the worker productivity, and all firms observe the same public signal. Therefore, due to Bertrand competition, wages are set by the expected productivity of workers. Specifically, let v, the productivity of workers from group g, be distributed normally with mean α_g and variance $\sigma_{v,g}$. Firms cannot observe the worker productivity, but they have access to a public noisy signal

$$y = v + u$$

2. To see this, notice that both firms make a wage offer uniformly on $[0, 0.5]$, and the winning wage offer is distributed with the CDF $F(x) = (\frac{w}{0.5})^2$ and the PDF $f(w) = 8w$. Therefore, the observed average workers wage is

$$\int_0^{0.5} w \times 8w dw = \frac{1}{3}$$

where the noise distributed normally $u \sim \mathcal{N}(0, \sigma_{u,g})$. Given the signal, the expected value of a worker's productivity is given by

$$E[v|y] = (1-\gamma)\alpha_g + \gamma y$$

where $\gamma = \frac{\sigma_{v,g}}{\sigma_{v,g}+\sigma_{u,g}}$. As discussed in Phelps (1972) and Aigner and Cain (1977), the resulting wage distribution for two groups of workers would be different if either the noisy signal or the underlying productivity are distributed differently across different groups. Specifically, we can see that as employers get a more precise signal, they will put higher weight on the signal in determining the wage, and rely less on the group mean. As Aigner and Cain (1977) note, with risk-neutral firms, the Phelps model implies that differences in average wages can only be explained by differences in workers' average productivity, which implies that in this statistical discrimination model, information is not enough to induce the observed wage gaps between groups and we need to assume that there exist differences in the underlying productivity distribution to rationalize the observed wage gaps.

As it turns out, this observation is more general than in the case of the normal distribution. Under the assumption that the market is competitive, and that firms are risk neutral, the differences in mean wages must be driven by differences in the underlying distribution, and cannot be explained by public signals, as shown in the claim below

Claim 5. Let g_1 and g_2 be two groups of workers. Assume that firms are risk neutral and observe the worker's group membership and a public signal $t_{g_i} \in \mathcal{T}_{g_i}$, drawn from a conditional distribution $\pi(t|v, g_i) \in \Delta(\mathcal{T})$. Assume that the observed mean wages of workers from group 1 and 2 are different, $\bar{w}_{g_1} \neq \bar{w}_{g_1}$, then it must be the case that $E[v|g_1] \neq E[v|g_2]$

Proof. First, notice that as all firms observe the same signal and are competing for the same worker, they are engaging in a Bertrand competition. As firms are risk neutral, this implies that all firms offer a wage that is equal to the expected value $w(t) = E[v|t, g_i]$. Then, notice

$$E[v|g_i] = E[E[v|t, g_i]|g_i] = E[w|g_i] = \bar{w}_{g_i}$$

Which implies that $\bar{w}_{g_1} \neq \bar{w}_{g_2} \implies E[v|g_1] \neq E[v|g_2]$ $\square$

As we know that averages of wages across gender and race are different, we know that in the setup shown in our model, it is not enough to assume that there is a set of signals, available to all firms, that can explain wage gaps, while holding the underlying productivity distributions the same across groups. Therefore, to examine the potential importance of information in explaining the wage gaps, we make the relaxation in our model that different firms may observe different signals on workers. This introduces an additional component to the strategic wage setting. Namely, firms need to make a guess on the worker's outside option. This additional strategic consideration can create a divergent between workers' productivity and their marginal output and as a result, generate wide wage gaps across groups with otherwise identical productivity distributions.

2.3 Partial Identification of Productivity Distribution and Inference

Our objective in this paper is to examine how much of the observed wage gap between different groups can be explained by differences in information access. We therefore can ask whether there exists a single distribution of workers productivity, $\mu \in \Delta(\mathcal{V})$, that can induce the observed wage gaps, with some information structure. More formally, let $H(w|g_i)$ be the observed wage cumulative distribution function (CDF) of workers from group i,[3] let $\kappa_k(w|t_k, g_i) = \int_0^w \beta_k(w|t_k, g_i)$ be the CDF of firm k wage offers, conditioned on the firms' signals. Finally, let $\kappa(w|t, g_i) = \prod_{k=1}^J \kappa_k(w|t_k, g_i)$ be

3. From here on, we suppress x for clarity

the predicted CDF for workers from group g, and some signal t. We want to examine whether the distribution of worker's productivity for workers of the two groups is the same. Specifically, we ask whether there exist two information structures and a distribution of workers productivity, such that, $\mu(v|g_1) = \mu(v|g_2) = \mu(v)$, and can generate the observed wage distributions, i.e.

$$H(w|g_i) = \int_{v,t} \kappa(w|t,g)\mathrm{P}(t|v,g)\mu(v)dvdt \tag{2.1}$$

As we are interested in the set of possible distributions μ that can generate the observed data, we now turn to explore how we can identify this set, within the basic model in section 2.2. First, throughout our analysis, we assume that the econometrician has access to data on wages, worker demographics and worker characteristics, such as education level or experience.

Assumption 6. *The econometrician observes the joint distribution $H(w,g,x)$, and their induced conditional probabilities. $H(w|g,x) \in \Delta(\mathcal{W})$.*

This assumption on the data available to the researcher is true for a large share of the empirical labor literature, which uses data on workers' wages but does not have access to data on workers' wage offers or productivity. Next, we define the set of model predictions to be the set of wage distributions that can result from the auction game with some information structure.[4]

Definition 2.3.1. The set of BNE predictions, $H \in \Delta(\mathcal{W})$, for a given information structure S and productivity distribution μ, is the of wage distribution induced by a BNE in the auction game

$$Q(S,\mu) = \{H : H(w) = \kappa(w|s)\mathrm{P}(s|v)\mu(v)\}$$

4. For clarity, we omit the group g indicator and add it when needed.

We also make the following assumption on the data generating process

Assumption 7. *The observed wage distribution is a result of a Bayes Nash Equilibrium in the labor-market auction game*

This assumption is quite strong, as the model we consider here is fairly restrictive. It does not allow workers to choose where to work based on job characteristics, other than wage. The model also assumes that all firms are homogeneous in their production technology and can extract the same output from workers. Although these are restrictive assumptions, they stress how - in an economy with almost no firm heterogeneity - information differences alone can generate a wide range of diverse outcomes. Finally, we can define the identified set of workers productivity distribution as

$$Q^{BNE}(H) = \{\mu : \exists S \in \mathcal{S} \text{ such that } H(w) \in Q(S.\mu)\}$$

This definition of the identified set may not seem useful because we need to iterate over all productivity distributions in $\Delta(\mathcal{V})$. For each distribution, we must find an information structure that induces the observed wage distribution. However, seminal results by Bergemann and Morris (Bergemann and Morris (2013), Bergemann and Morris (2016), Bergemann and Morris (2019b)) in information design and non-parametric estimation provide methods that transform this into a computationally feasible problem.

Before jumping to the result, it is worth introducing some notation. A game-form is a tuple $G = (\mathcal{W}, \mu)$ of the possible actions and prior distribution over the workers productivity. We define the a game to be the pair $(G, \mathcal{S})$.

Definition 2.3.2 (Bayes Correlated Equilibrium)**.** A joint distribution $\pi \in \Delta(\mathcal{V} \times \mathcal{W})$ is a Bayes Correlated Equilibrium of the basic form game $\mathcal{G}$, if for each firm j and wage offer w_j and deviation

w_j' we have

$$\sum_{v}\sum_{w_{-j}}\left[(v-w_j)q(w_j,\boldsymbol{w_{-j}})-S(v-w_j')q(w_j',\boldsymbol{w_{-j}})\right]\pi(v,w_j,\boldsymbol{w}_{-j})\geq 0 \quad \text{(Obedience Constraint)}$$

and the marginal of π with respect to the states is preserved

$$\sum_{w\in\mathcal{W}}\pi(v,\boldsymbol{w})=\mu(v) \quad \text{(prior consistency)}$$

Bergemann and Morris, shows that the set of distribution of actions and states, $\pi\in\Delta(v,\boldsymbol{w})$, that can be induced by BNE of $(\mathcal{G},\mathcal{S}')$, under some information structure $\mathcal{S}'$, is equivalent to the set of Bayes Correlated Equilibrium (BCE).

Theorem 1 (Bergemann and Morris (2016))**.** A distribution $\pi\in\Delta(\mathcal{V}\times\mathcal{W})$ that can arise as an outcome of a Bayes-Nash Equilibrium, under some information structure $\mathcal{S}$, if and only if it is a Bayes Correlated Equilibrium of the basic game $\mathcal{G}$

Next, we define the set of BCEs that can induce the observed wage distribution. Let π be a BCE, and let

$$BCE(H)=\left\{\pi:\sum_{\max(\boldsymbol{w})\leq w}\sum_{v}\pi(v,\boldsymbol{w})=H(w)\right\}$$

Similarly, we define set of productivity distributions, implied from the BCE, as the set of marginals over v

$$Q^{BCE}(H)=\{\mu:\pi\in BCE(H),\sum_{\boldsymbol{w}}\pi(v,\boldsymbol{w})=\mu(v)\}$$

Using Theorem 1, we have the following proposition

Proposition 7. The set of productivity distributions from a Bayes-Nash Equilibrium in the auction game is equal to the set of productivity distributions from a Bayes-Correlated Equilibrium in the basic game:

$$Q^{BCE} = Q^{BNE}$$

Proof. The proof follows trivially from the fact that that the set of BCEs is a convex set and Theorem 1. □

Therefore, proposition 7 shows us that it's enough to look for all the joint distributions of wage offers and workers productivity that can induce the observed wage distribution and satisfy the obedience and prior consistency constraints. In Appendix 2.8.2 we provide an illustrative example to show the identifying power of BCE in the case of one bidder.

2.3.1 Testing for the potential distorting effect of informational frictions

As discussed in the previous section, we want to see how much of the differences in the wage distribution can be attributed to information frictions. Following our discussion above, we can test whether a distribution μ can induce the observed wage distribution, with some information structure, by examining all the joint distributions π that have a marginal μ and satisfy the following constraints For every g we have

$$\begin{aligned}
\forall j, w, w' : \sum_{\boldsymbol{w}_{-j}, v} \pi(v, \boldsymbol{w}|g)\Big[(v-w)q(w, \boldsymbol{w_{-j}}) - (v-w')q(w, \boldsymbol{w_{-j}})\Big] \geq 0 \quad \text{(Obedience)} \\
\forall j : \sum_{v} \sum_{\boldsymbol{w}: w = max(\boldsymbol{w})} \pi(v, \boldsymbol{w}|g) = h(w) \quad \text{(Data-Match)}
\end{aligned} \tag{2.2}$$

where $h(w)$ is the density function of H. The first constraint is the obedience constraint, which, together with the third constraint, assures us that the resulting joint distribution of actions and states is a BCE, and therefore, there exists some BNE, with some information structure, that can induce it. The data match constraint, makes sure that the BCEs we consider can induce the observed wage distributions in the data.

As we are interested in the extent in which information, and not other underlying differences across groups, drives the size of the wage gap, we can first check whether there exist $\pi(v, \boldsymbol{w}|g_1)$ and $\pi(v, \boldsymbol{w}|g_2)$, that satisfy the linear constraint in 2.2 and

$$\sum_{\boldsymbol{w}} \pi(v, \boldsymbol{w}|g_1) = \sum_{\boldsymbol{w}} \pi(v, \boldsymbol{w}|g_2) \quad \forall v \in \mathcal{V} \tag{2.3}$$

Finding $\pi(v, \boldsymbol{w}|g_1)$ and $\pi(v, \boldsymbol{w}|g_2)$ that satisfies (2.2) and (2.3) assures us that there exists a single distribution μ that, with some information structure, can induce the wage distributions of the two groups. If such a distribution exists, then we cannot rule out the possibility that the observed wage gap between the two groups is induced by differences in the information firms have before making a job offer. If we cannot find a distribution that satisfies 2.2 and 2.3, then the differences in wages across groups are not driven solely by information frictions, but must be driven also by differences in the underlying productivity distribution.

Further more, we can also quantify the potential distorting effect of informational frictions in the labor market by finding the distribution of workers' productivity, implied by π, satisfying (2.2) that has the smallest mean and compare it to the observed mean wage. This would give us an upper bound on the potential size of information in shaping the wage distribution. Specifically, we want to measure

$$\begin{aligned} &\max \sum_{v} v \sum_{\boldsymbol{w}} \pi(v, \boldsymbol{w}|g) - \sum_{w} w h(w|g) \\ &s.t \ (2.2) \end{aligned} \tag{2.4}$$

The size of 2.4 gives us a bound on how much wages can diverge from the workers productivity and how much rents firms can extract from workers by utilizing their information.

2.3.1.1 Relation to other measures on discrimination

In the economics, discrimination is broadly categorized into statistical and taste-based paradigms, grounded in seminal works by Arrow (1973) and Becker (1957). Statistical discrimination involves decision-makers, commonly employers, using observable characteristics—such as race or gender—as heuristic proxies for unobservable attributes like skill or reliability, thereby generating biased outcomes (Phelps, 1972). Taste-based discrimination, by contrast, is rooted in the decision-maker's intrinsic preferences or prejudices against certain groups (Becker, 1957). Although these forms of discrimination have disparate motivations, both yield equivalently adverse impacts on marginalized populations.

Within the framework of our model, discrimination is entirely subsumed into the productivity distribution, μ. Specifically, if employers possess disutility in hiring from disadvantaged groups, this will manifest as a shift left in the distribution of productivity ν. Our metric for evaluating disparity is designed to answer the following question: In a setting where firms only have access to group membership information—assuming this constraint is also known to be shared by other firms—would wage offers be identical for individuals belonging to different groups?

Consequently, the model serves as a diagnostic tool: if we were to conclusively rule out the existence of a common productivity distribution across groups, it would imply that firms are incorporating race or group identity in their decision-making processes. Conversely, the identification of a parameter μ in line with our assumptions would suggest that the observed disparities between social groups could be attributed, at least partially, to additional information that firms possess either about individual productivity or competitive firms.

2.4 Computation and Inference

The set of joint distributions that satisfy (2.2) gives a tractable way to characterize the identified set of productivity distributions. Unfortunately, the size of π, the joint distribution, grows exponentially with the number of firms making a wage offer to the worker. For example, for a grid of size 15 and 10 firms, we need to keep track of 15^{11} variables. Therefore, If we represent the joint distribution as a vector of floats we would need around 35Gb of memory, and if we want to solve for the test for 2.3, we would need to hold twice as much memory. This clearly makes an analysis for a large number of players infeasible. Instead, we can use certain characteristics of the auction setup in order to reduce the dimensions of the problem.

We start by defining the set of identified productivity means to be

$$M = \{m = E[v; \mu] : \mu \in Q^{BNE}(H)\}$$

In Appendix 2.8.1 we show that this set is convex. This implies that it's enough to identify $\max(M)$ and $\min(M)$ to describe this set. Next, we show that we can restrict attention to a set of bi-mass distributions, that puts a positive mass only on 0, the lower bound of the support of the wage distribution, and $\bar{w} = \max(\mathcal{W}_i)$, the upper bound.

Claim 6. Let $\mu \in Q^{BCE}(H)$, then there exists a $\tilde{\mu}$ with two mass points on 0 and $\bar{w}$ and mean $E[v; \tilde{\mu}] = E[v; \mu]$ such that $\tilde{\mu} \in Q^{BCE}(H)$

The proof of this claim, as all other claims in this section is in Appendix 2.8.1. Next, we show that in order to check whether there exists a single distribution that can induce the wage distributions of two groups, then, it is sufficient to only check whether the set of means that can generate the wage distribution of one group, M_{g_1}, intersects with the set of means that can generate the wage distribution of the second group, M_{g_2}.

Claim 7. Let M_{g_i}, $i \in \{1,2\}$ be the set of identified means that can induce the wage distribution of group g_i. Then, there exists a distribution of worker productivity μ such that $\mu \in Q^{BCE}(H_{g_i})$ for $i \in \{1,2\}$ if and only if $M_{g_1} \cap M_{g_2} \neq \emptyset$. Also, the set of distribution means in $Q^{BCE}(H_{g_1}) \cap Q^{BCE}(H_{g_2})$ is contained in $[\max\{\underline{m}_{g_1}, \underline{m}_{g_2}\}, \min\{\overline{m}_{g_1}, \overline{m}_{g_2}\}]$ where $\overline{m}_{g_i} = \max(M_{g_i})$ and $\underline{m}_{g_i} = \min(M_{g_i})$.

Claim 6 and 7 and the fact that M is convex, implies that instead of characterising the entire set of possible distributions, we can just focus on finding M_{g_1} and M_{g_2} while restricting our search to a family of bi-mass distributions. This reduces the computational burden by, first, reducing the size of the joint distribution we need keep track of, and second, it allows us to solve the linear problem separately for each group and compare the set of identified means instead of solving the two problems together and require that (2.3) hold.

Finally, notice that the problem is grown exponentially with the number of players. We therefore want to solve a smaller problem, that takes advantage of our setup. We do it by taking advantage of the anonymous game structure of the auction game, and noticing that firms only care about the productivity of the worker, the highest wage offer, and the second highest wage offer. To take advantage of this we start by defining the object $p(w, w^1, n^1, w^2, n^2, v)$ which is the joint probability of a firm making a wage offer w, while the highest wage offer is w^1, the number of people who bid w^1 is n^1, the second highest wage offer is w^2 and n^2 is the number of firms who bid w^2. Notice that $p(w, w^1, n^1, w^2, n^2, v)$ has all the information needed to calculate the obedience and data match constraints. [5]

Furthermore, $p(w, w^1, n^1, w^2, n^2, v)$ does not grow exponentially with the number of players and therefore it is easier to work with, for larger set of players. We therefore, want to show that we can express the set of $Q^{BCE}(H)$ in terms of this object.

To do so, we start by requiring that $p(w, w^1, n^1, w^2, n^2, v)$ satisfy the obedience constraint

$$\sum_{w^1,n^1,w^2,n^2,v} p(w, w^1, n^1, w^2, n^2, v)((v-w)q(w, w^1, n^1, w^2, n^2)-(v-w')q(w', w'^1, n'^1, w'^2, n'^2)) \geq 0 \quad \forall w, w' \tag{2.6}$$

where (w'^1, n'^1, w'^2, n'^2) is the first and second order statistics of the modified distribution, if a firm changes it's action from w to w'. We also require that it satisfy the data match constraint

$$\sum_{w,n^1,w^2,n^2,v} p(w, \tilde{w}, n^1, w^2, n^2, v) = h(\tilde{w}) \tag{2.7}$$

The next set pf constraints assures that we have enough players to play against w^1 and w^2, in a

5. Notice that by defining $p(w, w^1, n^1, w^2, n^2, v)$ to be a distribution over the order statistics, we have also impose the following trivial constraints

$$\begin{aligned}
&w^1 \geq w \\
&w^1 \geq w^2 \\
&\text{if } w^1 = w^2 \text{ then } n^1 = n^2 > 1 \\
&\text{if } w^1 > w^2 \text{ then } n^1 = 1, n^1 + n^2 \leq N \\
&\text{if } n^1 = n^2 = N \text{ then } w = w^1 \\
&\text{if } n^1 + n^2 = N \text{ then } w \in \{w^1, w^2\} \\
&\text{if } w^1 > w^2 \text{ then } w \notin [w^2, w^1]
\end{aligned} \tag{2.5}$$

symmetric BCE[6]. The first constraint considers the case in which $w^1 > w^2$ and $n^1 + n^2 = N$

$$\frac{p(w^1, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^1-1}\binom{N-n^1}{n^2}} = \frac{p(w^2, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^2-1}\binom{N-n^2}{n^1}} \tag{2.8}$$

when $w^1 > w^2$, and $n^1 + n^2 < N$ we also require

$$\frac{p(w^1, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^1-1}\binom{N-n^1}{n^2}} = \sum_{w<w^2} \frac{p(w, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^1}\binom{N-1-n^1}{n^2}} \tag{2.9}$$

Finally, when $w^1 = w^2$, and $n^1 = n^2 < N$, we require that

$$\frac{p(w^1, w^1, n^1, w^1, n^1, v)}{\binom{N-1}{n^1-1}} = \sum_{w<w^1} \frac{p(w, w^1, n^1, w^1, n^1, v)}{\binom{N-1}{n^1}} \tag{2.10}$$

Denote by $BCEM(H)$ the set of marginals $p(w, w^1, n^1, w^2, n^2, v)$ that satisfy the above constraints for a given observed wage distribution H

$$BCEM(H) = \left\{ p(w, w^1, n^1, w^2, n^2, v) : p(w, w^1, n^1, w^2, n^2, v) \text{ satisfies } (2.5) - (2.10) \right\}$$

and let $Q^{BCE_M}(H)$ be the implied set of productivity distributions

$$Q^{BCEM}(H) = \left\{ \mu \in \Delta(\mathcal{V}) : \sum_{w,w^1,n^1,w^2,n^2} p(w, w^1, n^1, w^2, n^2, v) = \mu(v) \text{ and } p \in BCEM(H) \right\}$$

The next claim shows that the set of productivity distributions implied by any marginal in $BCEM(H)$ is the same as the set of productivity distribution we that can be rationalize with a BCE.

6. Claim 10 in the appendix shows that we can symmetrize any BCE when we use data only on the winning bids

Claim 8. $Q^{BCE}(H) = Q^{BCEM}(H)$

Therefore, instead of traversing the space of $BCEs$, we can use the restricted space of $BCEM$. This space does not grow exponentially with the number of players; rather, it grows quadratically, subject to additional constraints.

2.4.1 Inference

The identification arguments presented above assumed that we know the wage distribution H. However, when doing empirical analysis, we actually observe a an $i.i.d$ sample from the joint distribution $H(w, x, g)$, and therefore the analysis should take into account the sample variation. To do so, we follow the inference method suggested by Fang et al. (2020) for inference on linear systems with known coefficients. In what follows, we briefly describe the statistical test.

Given a $i.i.d$ sample of wages $\{w\}_i^n$ with w distributed according to $P \in \boldsymbol{P}$ Fang et al. (2020) show how to test the following hypothesis

$$H_0 : P \in \mathbf{P}_0 \quad H_1 : P \in \mathbf{P} \setminus \mathbf{P}_0$$

where

$$\mathbf{P}_0 \equiv \{P \in \mathbf{P} : \beta(P) = Ax \text{ for some } x \geq 0\}$$

where $A \in \mathbb{R}^{p \times d}$, with p as the number of constraints and d is the number of variables.[7] Fang et al. (2020) shows that in order to test whether x satisfies the linear problem, we can use the test statistics

7. It is known that any linear program with inequality constraints can be turned into a linear problem in standard form, in which all the inequalities are be written as equalities, with added slack variables. In our implementation we rewrite the linear problem in section 2.4 in its standard form

T_n

$$T_n \equiv \max\left\{\sup_{s\in\hat{\mathcal{V}}_n^{\mathrm{e}}} \sqrt{n}\left\langle s, \hat{\beta}_n - A\hat{x}_n^{\star}\right\rangle, \sup_{s\in\tilde{\mathcal{V}}_n^{\mathrm{i}}} \sqrt{n}\left\langle A^{\dagger}s, \hat{x}_n^{\star}\right\rangle\right\}$$

where $\hat{\beta}_n$ is an estimator for $\beta(P)$, which in our case is the density of the wage distribution and x_n^* is $A^{\dagger}\hat{\beta}_n$, in which $A^{\dagger}$ is the Moore-Penrose pseudoinverse of A. The test statics checks two types of violation - the first is whether $\hat{\beta}_n$ is in the range of A and the second is whether there exists $x \geq 0$, that solves the linear system. Fang et al. (2020) show how to calculate the critical value of the test by bootstrapping the sample $\hat{\beta}$ and solving a linear program at each iteration. The critical value they derive depends on a tuning parameter λ. We choose λ using the data-driven method they suggest.[8]

2.5 Data and Results

2.5.1 Data

We use the American Community Survey (ACS) 2010 sample to construct the wage distributions. We restrict our sample to individuals between the ages 21 and 65, who are in the labor force and are employed in the private sector. We remove self employed workers and restrict attention to workers who work full-time. We also remove people who earn at the top 1%. Figure 2.2 plots the wage distributions for white men, white women, black men and black women, and table 2.1 shows some descriptive on the workers from different groups. It is quite apparent from both the figure and the table that the two distributions are very different, where the distributions of women and black workers are more concentrated at low values.

Finally, to solve the linear program in (2.3) we normalize we normalize the wage distribution to

8. It seems that different values of λ do not change the results by much

be between $[0, H \times \frac{2}{3}]$ and discretize the set of bids as $\{0, .., H \times 2/3\}$. We let $H = 15$ which implies that we allow the highest value worker to be 1.5 times the maximum wage in our sample ($240,000).

2.5.2 Results

2.5.2.1 Market Frictions

Figure 2.1 shows the upper and lower bound on the average productivity of workers in dollars, per year, by demographic group, and under the assumption that there N firms who make a job offer to the worker.[9] The figure shows that the upper bounds on the productivity of white men is higher than that for the other groups. The lower bound for all groups is given simply by the mean wage (since under complete information, the observed wage distribution is the productivity distribution).[10] The figure also shows that as the number of firms who are making a wage offer increases, and therefore, the competition among firms intensifies, the set of productivity distributions that can induce the wage distribution shrinks. Table 2.3 shows the difference between the upper and lower bounds for each group and under the assumption that there are N firms offering wage. Notice that as the lower bound is given by the mean wage, this table presents the results to (2.4) and gives information on the potential distorting effect of information. We can see that as the number of firms who compete for workers is smaller, the potential role information can play is larger. For example, the difference between the mean productivity of white men and their average wage can go up to $114,000, if each worker only receives two wage offers. On the other hand, if there are less search frictions in the economy and each worker receives 50 wage offers, then the average wage can differ from the average

9. The bounds are showing the upper bound and lower bound of the confidence interval construct as described in section 4.1 and were calculated from the value of $N = \{2, 3, 5, 7, 10, 20, 50\}$

10. The small decline in driven by sample noise and our inference method

productivity level by roughly \$80,000. These bounds are not very tight, as the 90% of the wage distribution for white men is \$96,000. But these bounds capture the large role information can have in shaping the wage distribution.

The shaded area in figure 2.1 describes the set of bi-mass productivity distribution that can explain all four wage distributions. As discussed above, the set of these distribution decreases as the number of firms increases. Therefore, we can conclude that without imposing additional assumptions on the information set of the firms, we can't rule out that information frictions alone can explain all of the wage gap in the data.

Next, we turn to make a set of assumptions on the information set of the agent. First, we assume that workers sort into occupations and that it is common knowledge among all firms what is each worker's occupation. Tables 2.6, 2.7, 2.8 show the bounds on the mean productivity across different occupations, for different groups. First, we can see that the set of possible mean productivity for each occupation is wide. For example, the productivity for white men working at management, business, science and arts occupations can generate, on average between \$76,525 and \$208,894 and on the other side, workers in production can generate on average between \$38,514 and \$133,732. Interestingly, we find that we cannot rule out that workers in all occupations have the same average productivity. In our setup, information can give rise to differences in workers' wage across occupations, even if the distribution of workers productivity is the same across all occupations. For example, firms might find it much harder to assess whether a worker is going to be a good manager or not than it would assessing whether a worker would do a good job in the assembly line. These differences in the available information to firms can generate the observed differences in wages, rather than self-selection of different quality workers or the role of each occupation in the production process.

Next, we impose the assumption that all firms observe the workers' experience and education level.[11] We divide the education level to three categories - High school dropouts, high-school graduates/have some college education, and workers with college degree. Similarly, we divide the experience level into three groups - 0-6, 6-12 and more than 12 years. This partition captures the shape of the wage schedule, as discussed in Rubinstein and Weiss (2006).

A common practice in the empirical labor literature is to condition wages on both experience and education. This goes back to Mincer (1958), who rationalized the linear structure of the wage equation using compensation differential arguments. Later papers justify the inclusion of these variables in wage equation based on a human capital rationale (Heckman et al. (2005)), implying that workers' ability changes as they acquire education and experience on-the-job training. In most of these models, workers are being compensated by firms, which are assumed to observe the investments workers are making in human capital. In the framework we present, this amounts to an assumption on the information available to the firms. Specifically, we assume that all firms observe the workers investment in education and the experienced they gained.

Table 2.4 and 2.5 estimate the bounds on the mean productivity, under the assumption that firms observe a public signal on the workers education level and experience. Interestingly, we find that for relatively low level of competition, wage disparities between highly educated and experienced white men and uneducated and inexperienced white men cannot be explained by a single productivity distribution and different signals observed by the firms. This implies that, under the assumption that all firms observe workers' education level and experience, workers' ability differs between experienced educated workers and non experienced educated workers. We again cannot rule out that there are no differences between the four groups of workers.

In our latest exercise, we examine whether all the information frictions are driven by different

11. Following the convention, we define experience to be $Age - 6 -$ School Years

selection patterns. In our model, selection and sorting into different industries can be thought of as components of the signal and information firms possess. For instance, when not conditioning on industry, part of the information firms hold may include the industry in which they operate, and the informational content of the industry is influenced by the varying selection patterns within these sectors. Consider a simplified model where both men and women are know their latent productivity. Further assume that women are discouraged from pursuing STEM fields before entering the labor market. In this scenario, the women who do choose STEM are likely to have higher latent abilities. In our framework, different market structures can be interpreted as signals. Consequently, we may wish to tighten the conditions of our test to exclude cases where selection is not informative, conditional on group membership. This would imply that $\mu(v|g, Industry) = \mu(v|Industry)$. In such a world, agents may choose industries differently, but these variances in selection are not stratified by group.

Table 2.9, 2.10, 2.11 present results on mean average productivity. Our findings suggest that we cannot dismiss the possibility that differences in average wages are influenced by factors other than selection. Thus, we demonstrate that information frictions can account for wage gaps even when selection patterns are consistent across industries.

To further strengthen the test, we investigate whether wage disparities can persist in the absence of selection across all industries and groups. For this, we require that $\mu(v|g, Industry) = \mu(v)$. In scenarios with two competing firms, the bounds on the mean wage distribution remain largely unchanged, falling between $(481266, 1318475)$, implying that other informational factors can continue to influence the observed wage gap.

2.6 Conclusion

In this paper, we explored the potential role information frictions play in shaping the wage gaps. We found that differences in average wages across white men, white women, black men and black women can be explained only by information frictions. This result differs from previous results in the statistical discrimination literature that argued that incomplete information is not enough to explain average wage gaps between groups of workers. We find that the simple model can generate the observed wage distribution without the need to argue for differences in the underlying productivity distribution of workers. This paper stresses the potential importance that information may have on the wage setting process. It implies that additional research is needed to understand what firms know about their job applicants and the applicants' outside option. Within the framework we use here, it will be interesting to explore further what assumptions we need to impose on the accuracy of the information firms have, to be certain that information frictions are not the sole reason for observed wage gaps. Also, leveraging the results from Bergemann et al. (2017) for the lowest possible revenue, over all information structures, we can try and see what is the largest wage gap possible that can be driven solely by differences in information. Finally, throughout the paper we make a strong assumption that firms know the number of competitive wage offers. Following Bergemann et al. (2021b) we can try and relax this assumption and see how this affects the set of identified distributions.

2.7 Tables and Figures

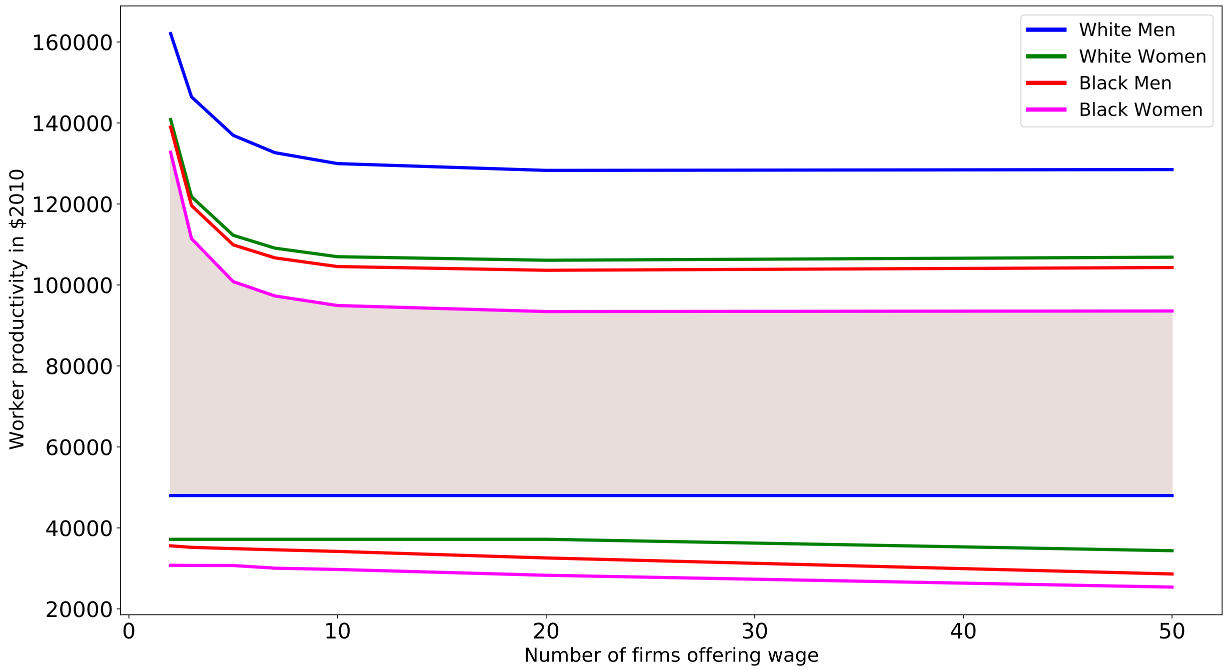

Figure 2.1: Upper and Lower bound on the average productivity of workers

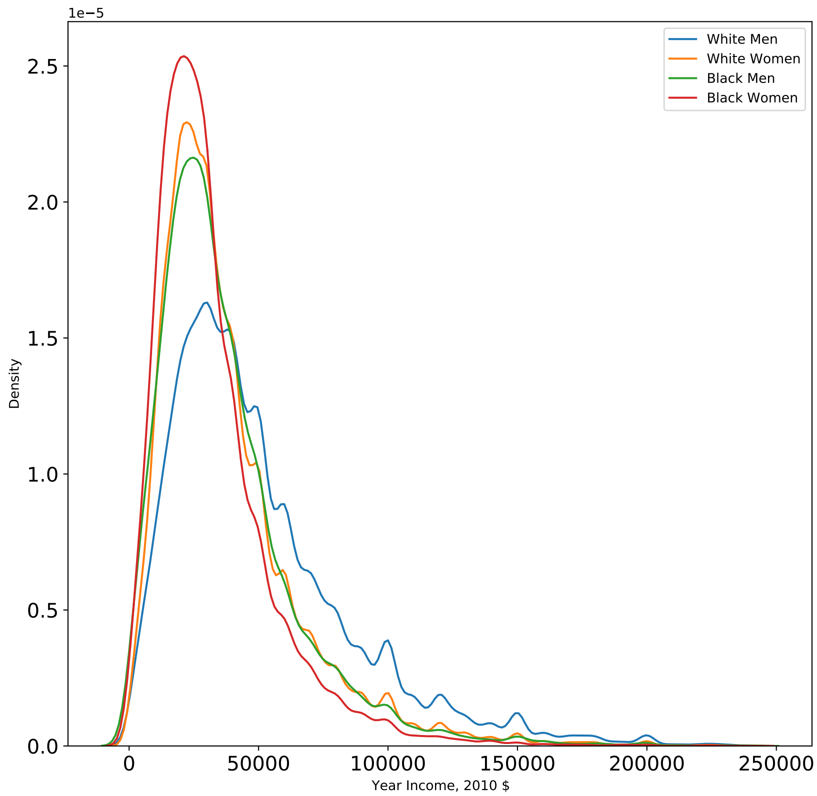

Figure 2.2: The four groups wage density

	WM	WW	BM	BW
Mean	48563.17	37554.15	36432.41	31341.16
Max	240000	240000	240000	240000
Min	100	10	160	270
5%	10000	8600	6011	6400
10%	14400	12000	11000	10000
25%	24000	20000	19200	16800
50%	40000	30000	30000	26000
75%	63000	48000	46996	40000
90%	96000	71000	70000	60000
95%	120000	90000	86000	74000

Table 2.1: Descriptive Statistics

Number of firms making wage offer	WM	WW	BM	BW
2	[47987,162073]	[37179,140853]	[35573,138985]	[30766,132775]
3	[47988,146452]	[37188,121794]	[35199,119605]	[30718,111444]
5	[47988,136953]	[37188,112263]	[34881,109901]	[30707,100821]
7	[47988,132665]	[37188,109088]	[34596,106686]	[30056,97281]
10	[47988,129968]	[37188,106971]	[34193,104542]	[29736,94920]
20	[47988,128291]	[37188,106100]	[32577,103624]	[28301,93431]
50	[47988,128493]	[34368,106859]	[28620,104303]	[25384,93569]

Table 2.2: The potential effect of information frictions - Lower and upper bound on mean productivity

Number of firms making a wage offer						
2	3	5	7	10	20	50
[47987,132775]	[47988,111444]	[47988,100821]	[47988,97281]	[47988,94920]	[47988,93431]	[47988,93569]

Table 2.3: The potential effect of information frictions - Lower and upper bound on the mean productivity of distribution who can explain the four wage distributions

		2 Firms			
Education Level	Experience	WM	WW	BM	BW
High School Dropout	0-6	[15319,92703]	[11081,87987]	[9011,96062]	[7092,94718]
	7-12	[20342,96624]	[16152,94052]	[16435,103435]	[14517,94652]
	> 12	[28838,122783]	[20210,92011]	[24767,117984]	[19899,95538]
High School Graduate	0-6	[21256,99727]	[18746,90545]	[17669,93798]	[16549,92353]
	7-12	[31603,126274]	[26236,111052]	[25679,119453]	[22478,102406]
	> 12	[45527,148205]	[34188,129723]	[35864,133170]	[29737,125400]
Colledge Degree	0-6	[40168,144552]	[35318,131746]	[31989,136923]	[30233,129410]
	7-12	[62895,181397]	[52576,164796]	[45543,158272]	[45392,148662]
	> 12	[82479,212782]	[62469,184971]	[63189,190784]	[54328,171938]
		7 Firms			
High School Dropout	0-6	[14029,66141]	[10583,57881]	[7281,65979]	[6608,62251]
	7-12	[19533,73363]	[15407,67793]	[14517,79088]	[11659,67656]
	> 12	[27643,89631]	[19345,68390]	[22643,89364]	[18391,71701]
High School Graduate	0-6	[21025,76537]	[18745,65859]	[17669,68383]	[15921,66302]
	7-12	[30797,92048]	[25448,84638]	[24225,88884]	[21381,80770]
	> 12	[45530,120607]	[34196,96999]	[34821,101871]	[28827,89957]
Colledge Degree	0-6	[38983,116040]	[34111,100075]	[28947,104308]	[28086,95252]
	7-12	[62879,148784]	[52584,130671]	[41184,126222]	[42260,123713]
	> 12	[82479,176048]	[62481,151496]	[61468,155905]	[51629,135734]

Table 2.4: Bounds on workers average productivity, conditional on education level and potential experience

		10 Firms			
Education Level	Experience	WM	WW	BM	BW
High School Dropout	0-6	[13402,64374]	[9210,55929]	[5595,63972]	[6352,60087]
	7-12	[18774,71814]	[14893,66074]	[13406,77559]	[10829,65936]
	> 12	[27146,87215]	[18788,66823]	[21884,86833]	[17799,70146]
High School Graduate	0-6	[20880,74991]	[18745,64213]	[17669,66683]	[15599,64571]
	7-12	[30210,89768]	[25181,82102]	[23488,86375]	[20952,79312]
	> 12	[45530,117837]	[34196,94823]	[34379,99804]	[28420,87588]
Colledge Degree	0-6	[37989,114158]	[33622,97962]	[27790,102131]	[27246,92988]
	7-12	[60372,145800]	[52582,128350]	[38531,123631]	[40364,120912]
	> 12	[82479,174027]	[62481,148403]	[58411,153232]	[50631,133584]
		20 Firms			
High School Dropout	0-6	[11840,63242]	[7294,54182]	[3492,62351]	[4524,58341]
	7-12	[17532,71316]	[12464,65073]	[9429,76783]	[7855,64738]
	> 12	[26031,85318]	[17814,65963]	[18931,84802]	[16250,69312]
High School Graduate	0-6	[19643,74645]	[17464,63243]	[14190,65654]	[13672,63339]
	7-12	[29067,88216]	[24222,79797]	[21368,84287]	[19135,78954]
	> 12	[45530,115852]	[34197,93631]	[33062,98990]	[27361,85758]
Colledge Degree	0-6	[35849,112219]	[32144,97114]	[23457,101523]	[21664,91743]
	7-12	[56893,143332]	[52569,127208]	[31380,121825]	[35756,118823]
	> 12	[82479,171734]	[62481,145958]	[52517,150815]	[45987,132912]

Table 2.5: Bounds on workers average productivity, conditional on education level and potential experience

	2 Firms			
Occupation Group	WM	WW	BM	BW
Management, Business, Science, and Arts Occupations	[76625,208894]	[57573,181048]	[58336,187206]	[48810,168396]
Business Operations Specialists	[65573,186839]	[52541,164298]	[48725,176692]	[43796,151843]
Financial Specialists	[71936,199559]	[52686,162749]	[52865,178582]	[43615,143603]
Computer and Mathematical Occupations	[74141,196949]	[63720,181718]	[60328,179537]	[52433,166268]
Architecture and Engineering Occupations	[71872,192156]	[55976,172515]	[62691,185354]	[53822,165105]
Life, Physical, and Social Science Occupations	[67111,196682]	[52802,175720]	[44551,171823]	[41623,154256]
Community and Social Services Occupations	[37646,138424]	[37043,130940]	[33372,132875]	[31944,125830]
Legal Occupations	[88543,235467]	[56902,174134]	[60966,225476]	[48260,181602]
Education, Training, and Library Occupations	[47662,163162]	[32599,132008]	[37783,157180]	[28892,132628]
Arts, Design, Entertainment, Sports, and Media Occupations	[53833,176646]	[44932,154469]	[44734,166339]	[41846,176973]
Healthcare Practitioners and Technical Occupations	[66053,197836]	[49731,156613]	[52702,179129]	[45163,153952]
Healthcare Support Occupations	[27874,130456]	[24848,100183]	[24823,121753]	[23327,100108]
Protective Service Occupations	[32309,139368]	[28109,131308]	[28286,133987]	[24715,120371]
Food Preparation and Serving Occupations	[20828,97493]	[17492,88966]	[19136,98924]	[16585,87479]
Building and Grounds Cleaning and Maintenance Occupations	[26093,116161]	[18498,89078]	[22015,105658]	[17688,89980]
Personal Care and Service Occupations	[30939,132934]	[21914,103643]	[23918,119923]	[20125,97484]
Sales and Related Occupations	[50945,168246]	[33413,140368]	[35719,143231]	[23720,115050]
Office and Administrative Support Occupations	[36085,136175]	[31668,123477]	[29733,127330]	[28946,122347]
Farming, Fishing, and Forestry Occupations	[27890,127262]	[39127,152968]	[25718,149189]	[39177,153329]
Construction and Extraction Occupations	[37479,138066]	[30028,136089]	[32630,135689]	[24688,144758]
Extraction Workers	[44358,157283]	[21782,152090]	[26276,163722]	NA
Installation, Maintenance, and Repair Workers	[42828,139650]	[37730,142002]	[37729,136710]	NA
Production Occupations	[38514,133732]	[26918,111167]	[32992,130707]	NA
Transportation and Material Moving Occupations	[33802,131167]	[24511,111238]	[29967,127651]	NA

Table 2.6: Bounds on workers average productivity, conditional on workers occupation

	7 Firms			
Occupation Group	WM	WW	BM	BW
Management, Business, Science, and Arts Occupations	[76625,169777]	[57577,145076]	[54116,153555]	[45329,135079]
Business Operations Specialists	[65448,154220]	[52099,131696]	[42098,138651]	[39345,122227]
Financial Specialists	[71918,163953]	[52485,130089]	[48050,142737]	[39945,117936]
Computer and Mathematical Occupations	[74051,156648]	[61596,146183]	[56233,147102]	[47310,131507]
Architecture and Engineering Occupations	[70479,154077]	[54610,138489]	[55688,150741]	[46836,132635]
Life, Physical, and Social Science Occupations	[65720,160237]	[51258,138811]	[39839,133847]	[37862,119501]
Community and Social Services Occupations	[33007,109672]	[34161,101851]	[26580,104162]	[29032,93117]
Legal Occupations	[88496,200619]	[56878,144154]	[53748,190271]	[38300,145431]
Education, Training, and Library Occupations	[45377,127190]	[31392,98486]	[33176,124857]	[26541,97419]
Arts, Design, Entertainment, Sports, and Media Occupations	[52800,139921]	[44850,127186]	[36617,131619]	[36937,138408]
Healthcare Practitioners and Technical Occupations	[65886,162302]	[48325,123964]	[45033,150164]	[42621,122434]
Healthcare Support Occupations	[24961,98421]	[23822,79817]	[22091,88241]	[22076,78861]
Protective Service Occupations	[29757,105633]	[24337,97460]	[26597,99728]	[21001,93328]
Food Preparation and Serving Occupations	[19804,74132]	[17088,63444]	[17599,75130]	[15758,61597]
Building and Grounds Cleaning and Maintenance Occupations	[25176,88096]	[17906,64631]	[20791,83971]	[16728,65258]
Personal Care and Service Occupations	[28405,98663]	[20846,81088]	[22266,85450]	[19376,73986]
Sales and Related Occupations	[50911,138401]	[32574,105960]	[32731,111553]	[21909,91027]
Office and Administrative Support Occupations	[36063,104420]	[30972,89282]	[27728,91958]	[27952,86430]
Farming, Fishing, and Forestry Occupations	[25201,99290]	[35471,123528]	[21461,112086]	[30778,126935]
Construction and Extraction Occupations	[36623,107669]	[26634,101684]	[29600,102368]	[24274,108122]
Extraction Workers	[40273,125597]	[18606,120786]	[14447,119779]	NA
Installation, Maintenance, and Repair Workers	[42799,113402]	[37258,113339]	[35077,107657]	NA
Production Occupations	[38518,104185]	[25519,88313]	[30712,97642]	NA
Transportation and Material Moving Occupations	[33793,98113]	[23298,86653]	[28609,92708]	NA

Table 2.7: Bounds on workers average productivity, conditional on workers occupation

	20 Firms			
Occupation Group	WM	WW	BM	BW
Management, Business, Science, and Arts Occupations	[76625,166430]	[57577,141935]	[45345,148821]	[37079,131489]
Business Operations Specialists	[58736,149263]	[46327,128267]	[29731,135836]	[31570,117638]
Financial Specialists	[71885,159846]	[45646,126686]	[33171,140255]	[33936,112715]
Computer and Mathematical Occupations	[68836,153449]	[56905,141257]	[46843,141379]	[38030,129238]
Architecture and Engineering Occupations	[67924,150277]	[46086,136264]	[45065,146401]	[34522,131654]
Life, Physical, and Social Science Occupations	[55480,155727]	[40652,135961]	[28514,130594]	[21564,115947]
Community and Social Services Occupations	[24490,107272]	[28945,99429]	[16578,102073]	[22298,90054]
Legal Occupations	[88496,195748]	[47229,140724]	[53748,184538]	[31877,142224]
Education, Training, and Library Occupations	[38981,123319]	[28460,95193]	[17513,120517]	[22330,94152]
Arts, Design, Entertainment, Sports, and Media Occupations	[47738,136981]	[37656,122293]	[24753,128216]	[26116,135459]
Healthcare Practitioners and Technical Occupations	[58297,157688]	[45870,119938]	[32726,146931]	[36208,117997]
Healthcare Support Occupations	[20014,94702]	[22574,78217]	[16147,83810]	[19959,77462]
Protective Service Occupations	[24897,102426]	[18697,93604]	[20706,95901]	[15724,88592]
Food Preparation and Serving Occupations	[18077,71971]	[15925,60509]	[15409,73057]	[14335,58548]
Building and Grounds Cleaning and Maintenance Occupations	[23477,83498]	[16375,61965]	[18349,78770]	[14886,62529]
Personal Care and Service Occupations	[24228,95010]	[19357,79344]	[19014,80651]	[17723,71946]
Sales and Related Occupations	[47882,134475]	[30151,102301]	[25232,108178]	[18490,85934]
Office and Administrative Support Occupations	[36060,101409]	[29988,85300]	[24901,88066]	[26061,82214]
Farming, Fishing, and Forestry Occupations	[20928,94898]	[30530,118689]	[11645,108258]	[18845,122541]
Construction and Extraction Occupations	[34674,105381]	[18908,98501]	[24187,99302]	[8308,101214]
Extraction Workers	[31598,120985]	[11504,117622]	[12112,114695]	NA
Installation, Maintenance, and Repair Workers	[39817,108685]	[26796,109516]	[30595,105905]	NA
Production Occupations	[35848,101749]	[23603,83685]	[27117,94170]	NA
Transportation and Material Moving Occupations	[30841,94733]	[20671,81744]	[25467,88790]	NA

Table 2.8: Bounds on workers average productivity, conditional on workers occupation

	2 Firms			
Industry	WM	WW	BM	BW
Manufacturing	[51603,165744]	[41228,147681]	[37974,140989]	[32236,135136]
Agriculture, Forestry, Fishing and Hunting	[26349,121153]	[20759,114528]	[20276,114465]	[13559,117558]
Mining, Quarrying, and Oil and Gas Extraction	[58573,182287]	[49893,174495]	[42515,171245]	[43006,180193]
Utilities	[68786,188346]	[55564,171634]	[52700,179344]	[45394,164217]
Construction	[41245,145932]	[38722,136342]	[32555,133411]	[34464,161730]
Wholesale Trade	[50023,161877]	[40677,146380]	[34408,134032]	[31644,139270]
Retail Trade	[38702,144191]	[29245,121394]	[29409,132853]	[23726,109194]
Transportation and Warehousing	[44925,151051]	[34796,130588]	[35069,135860]	[31672,128256]
Information	[59357,182527]	[48665,163438]	[47853,164229]	[40560,145176]
Finance and Insurance	[68880,198556]	[45656,150463]	[48799,168783]	[38032,136619]
Real Estate and Rental and Leasing	[45301,157944]	[38286,141775]	[31590,135128]	[28701,126959]
Professional, Scientific, and Technical Services	[73554,204013]	[50989,163977]	[54941,182658]	[45757,160898]
Management of Companies and Enterprises	[76108,221525]	[50109,178051]	[38003,209860]	[30729,161417]
Administrative and Support and Waste Management and Remediation Services	[34711,142735]	[30748,135058]	[27053,127743]	[25166,119261]
Educational Services	[44992,155291]	[34927,132773]	[35322,143380]	[33985,132721]
Health Care and Social Assistance	[50806,170414]	[36929,137531]	[34237,138185]	[30027,129365]
Arts, Entertainment, and Recreation	[34991,142072]	[27996,126447]	[26216,129311]	[22568,116629]
Accommodation and Food Services	[25733,115601]	[20555,100725]	[22556,113567]	[17749,97714]
Other Services (except Public Administration)	[35504,134105]	[23456,108539]	[27857,132699]	[20734,99251]

Table 2.9: Bounds on workers average productivity, conditional on workers industry

	7 Firms			
Industry	WM	WW	BM	BW
Manufacturing	[51612,134608]	[41231,118723]	[36726,110687]	[30235,101586]
Agriculture, Forestry, Fishing and Hunting	[24607,92681]	[18941,90768]	[17736,89932]	[8749,93443]
Mining, Quarrying, and Oil and Gas Extraction	[55265,148578]	[45856,141418]	[32556,129821]	[27710,140728]
Utilities	[66263,149144]	[51783,136749]	[48326,145166]	[38776,126950]
Construction	[41239,117385]	[35609,107135]	[29524,100374]	[29389,127405]
Wholesale Trade	[48910,133800]	[38893,117560]	[31052,102648]	[25541,105761]
Retail Trade	[38710,113214]	[28570,95350]	[27410,97872]	[22440,87595]
Transportation and Warehousing	[44052,123077]	[32990,98875]	[32898,105019]	[29800,95238]
Information	[59334,148093]	[46611,132647]	[43232,127500]	[35926,117606]
Finance and Insurance	[68880,162241]	[45670,124398]	[42958,136879]	[35404,107319]
Real Estate and Rental and Leasing	[43829,131150]	[35568,111745]	[25634,100984]	[24940,91391]
Professional, Scientific, and Technical Services	[73558,165286]	[50963,133384]	[52721,146898]	[42123,128791]
Management of Companies and Enterprises	[68437,184609]	[44744,140143]	[38003,188899]	[18735,114030]
Administrative and Support and Waste Management and Remediation Services	[33106,109539]	[29329,99664]	[24243,96567]	[23813,90303]
Educational Services	[43031,125904]	[33431,101282]	[32062,114458]	[30621,101375]
Health Care and Social Assistance	[50829,139442]	[36207,106364]	[30652,105565]	[28557,96175]
Arts, Entertainment, and Recreation	[34464,109715]	[25202,94023]	[22697,91891]	[19824,88929]
Accommodation and Food Services	[24510,92875]	[19695,76858]	[20192,90335]	[16580,72427]
Other Services (except Public Administration)	[33874,102865]	[22392,86758]	[24614,97725]	[19042,76866]

Table 2.10: Bounds on workers average productivity, conditional on workers industry

	20 Firms			
Industry	WM	WW	BM	BW
Manufacturing	[51612,130834]	[37980,116419]	[32591,108308]	[26179,98146]
Agriculture, Forestry, Fishing and Hunting	[21109,88130]	[14737,85484]	[13076,84088]	[3480,91745]
Mining, Quarrying, and Oil and Gas Extraction	[48265,145092]	[33985,138430]	[24853,127153]	[27710,134493]
Utilities	[61672,144757]	[44739,134383]	[33646,139736]	[28042,122922]
Construction	[38744,114453]	[31312,104486]	[24550,97369]	[17793,123590]
Wholesale Trade	[45915,129799]	[35197,115140]	[24300,99105]	[17443,102681]
Retail Trade	[36154,110269]	[26704,90931]	[23412,93794]	[18857,84367]
Transportation and Warehousing	[41544,118305]	[29460,95778]	[28707,102418]	[25914,92120]
Information	[53656,144592]	NA	[34904,123931]	NA
Finance and Insurance	[68880,157830]	[42065,119983]	[32769,133746]	[30948,105078]
Real Estate and Rental and Leasing	[37659,127280]	[31054,109052]	[18259,97315]	[18567,87449]
Professional, Scientific, and Technical Services	[73558,161522]	[46587,129410]	[41286,143796]	[35427,124724]
Management of Companies and Enterprises	[57786,180154]	[30070,137348]	[38003,187760]	[16930,106873]
Administrative and Support and Waste Management and Remediation Services	[29293,106325]	[25667,95805]	[19796,92246]	[19712,85836]
Educational Services	[36777,121483]	[30639,98317]	[23872,110801]	[25647,98710]
Health Care and Social Assistance	[43814,135558]	[34379,103685]	[24682,102088]	[26470,92274]
Arts, Entertainment, and Recreation	[26851,106333]	[21969,89930]	[16239,87826]	[14144,83966]
Accommodation and Food Services	[22004,87941]	[17967,74770]	[17136,85375]	[13866,69803]
Other Services (except Public Administration)	[31306,100058]	[19250,83238]	[17398,93708]	[16398,75103]

Table 2.11: Bounds on workers productivity, conditional on workers industry

2.8 Appendix

2.8.1 Proofs

Claim 9. The set of identified means,

$$M = \{m = E[v; \mu] : \mu \in Q^{BNE}(H)\}$$

is convex.

Proof. fix $m^*, m^{**} \in M$ and choose $m \in [m^*, m^{**}]$ and λ such that $\lambda m^* + (1-\lambda)m^{**} = m$. We want to show that there exists a joint distribution π with marginal $\sum_w \pi(v,w) = \mu(v)$ and $E[v;\mu] = m$ such that μ is part of the identified set of distributions. Let π^* and π^{**} be two BCEs that induce H and have marginals μ^* and μ^{**} with the corresponding means. We can then define define π to be $\lambda\pi^*(v,\boldsymbol{w}) + (1-\lambda)\pi^{**}(v,\boldsymbol{w})$. Notice that for each v we have

$$\sum_w \pi(v,w) = \sum_w \lambda\pi^*(v,\boldsymbol{w}) + (1-\lambda)\pi^{**}(v,\boldsymbol{w}) = \mu(v)$$

Similarly, π satisfies the data match constraint

$$\begin{aligned}\sum_v \sum_{\boldsymbol{w}:\max(\boldsymbol{w})=w} \pi(v,\boldsymbol{w}) &= \sum_v \sum_{\boldsymbol{w}:\max(\boldsymbol{w})=w} \lambda\pi^*(v,\boldsymbol{w}) + (1-\lambda)\pi^{**}(v,\boldsymbol{w}) \\ &= \lambda H(w) + (1-\lambda)H(w) \\ &= H(w)\end{aligned}$$

and also the obedience constraint

$$\sum_v \sum_{w_{-j}} \pi(v,\boldsymbol{w})\Delta(w_j,w',w_{-j},v) = \sum_v \sum_{w_{-j}} \lambda\pi^*(v,\boldsymbol{w}) + (1-\lambda)\pi^*(v,\boldsymbol{w})\Delta(w_j,w',w_{-j},v) \geq 0$$

Therefore π is a BCE that induces the wage distribution H and $m \in M$ $\square$

2.8.1.1 Proof of Claim 6

Proof. Let $\mu \in Q^{BCE}(H)$ and fix a π such that $\sum_w \pi(v, \boldsymbol{w}) = \mu(v)$ and π induces H. Then notice

$$
\begin{aligned}
0 \leq \sum_{v, \boldsymbol{w}_{-i}} \pi(v, \boldsymbol{w}) \Big[(v - w_k) q(w_k, \boldsymbol{w_{-k}}) - (v - w_k') q(w_k, \boldsymbol{w_{-k}}) \Big] = \\
\sum_{v, \boldsymbol{w}_{-i}} p(\boldsymbol{w}) F(v|\boldsymbol{w}) \Big[v(q(w_k, \boldsymbol{w_{-k}}) - q(w_k, \boldsymbol{w_{-k}})) + (w_k q(w_k, \boldsymbol{w_{-k}}) - w_k' q(w_k', \boldsymbol{w_{-k}}) \Big] = \\
\sum_{\boldsymbol{w}_{-i}} p(\boldsymbol{w}) \Big[E[v|\boldsymbol{w}](q(w_k, \boldsymbol{w_{-k}}) - q(w_k, \boldsymbol{w_{-k}})) + (w_k q(w_k, \boldsymbol{w_{-k}}) - w_k' q(w_k', \boldsymbol{w_{-k}}) \Big]
\end{aligned}
$$

We can therefore construct the following $\tilde{\pi}$ by equating the marginals $\sum_v \pi(v, \boldsymbol{w}) = \sum_v \pi(v, \boldsymbol{w})$ and defining

$$
\begin{aligned}
\tilde{\pi}(\bar{w}|\boldsymbol{w}) &= p \text{ s.t } p \times \bar{w} = E[v|\boldsymbol{w}] \\
\tilde{\pi}(0|\boldsymbol{w}) &= 1 - \pi(\bar{w}|\boldsymbol{w}) \\
\forall v \notin \{0, H\}, \tilde{\pi}(v|\boldsymbol{w}) &= 0.
\end{aligned}
$$

Notice that by construction $\tilde{\pi}$ satisfies both the obedience constraint and data match constraint and therefore $\tilde{\pi} \in Q^{BCE}(H)$. Finaly, let $\tilde{\mu} = \sum_{\boldsymbol{w}} \tilde{\pi}(v, \boldsymbol{w})$, and notice that due to the law of iterated expectations we have that $E[v; \tilde{\mu}] = E[v; \mu]$ as needed. □

2.8.1.2 Proof of claim 7

Proof. The first direction is easy. If $M_{g_1} \cap M_{g_2} \neq \emptyset$ then we know that $Q^{BCE}(H_{g1}) \cap Q^{BCE}(H_{g2}) = \emptyset$. We prove the reverse direction by construction. Let $\mu_{g_1} \in Q^{BCE}(H_{g1}), \mu_{g_2} \in Q^{BCE}(H_{g2})$ and have the same mean $m_{g_1} = m_{g_2}$. We want to show that there exist at least one distribution that can rationalize both distributions. From claim 6, we know that we can construct a distribution

$\tilde{\mu}_{g_i} \in Q^{BCE}(H_{g_i})$, with two mass points on the edges of the support and $m_{g_i} = E[v; \tilde{m}]$. Therefore we can construct two such distributions $\tilde{\mu}_{g_1}$ and $\tilde{\mu}_{g_2}$. But as $E[v; \tilde{m u}_{g_1}] = E[v; \tilde{m u}_{g_2}]$, then it must be $\tilde{\mu}_{g_1} \overset{d}{=} \tilde{\mu}_{g_2}$, as needed. Further notice that we can do this for each mean value in the interval $[\max\{\underline{m}_{g_1}, \underline{m}_{g_2}\}, \min\{\overline{m}_{g_1}, \overline{m}_{g_2}\}]$, which concludes the proof. □

2.8.1.3 Proof of claim 8

Before proving claim **??**, we show that if we only have access to wages, and not wage offers, it is without loss to restrict attention only to a symmetric (i.e. exchangeable) BCEs

Claim 10. For any $\pi \in BCE(H)$, there exists a symmetrized $\tilde{\pi}$ that is also in $BCE(H)$.

Proof. We want to show that there exist and exchangable BCE $\tilde{\pi}(v, \boldsymbol{w})$ that can induce the same winning bid. We show this by construction. Let Ξ be the set of permutations of $\{1, ...N\}$ and we associate each permutation with a mapping from $\mathcal{W}^N \to \mathcal{W}^N$ where $\xi(\boldsymbol{w})$ is a permuted profile of wage offers, in which $\xi_i(\boldsymbol{w}) = w_{\xi(i)}$. First, notice that any permuation of the players in a BCE is also a BCE. Then, fix $\pi \in BCE(H)$, and define define $\tilde{\pi}$ to be

$$\tilde{\pi}(v, \boldsymbol{w}) = \frac{1}{N!} \sum_{\xi \in \Xi} \pi(v, \xi(\boldsymbol{w}))$$

and notice that $\tilde{\pi}$ satisfies the obedience constraint and the prior consistency constraint and therefore a BCE. Further notice that it can generate the winning bid distribution

$$\begin{aligned}
\sum_{v} \sum_{\boldsymbol{w}:\max(\boldsymbol{w})=w} \tilde{\pi}(v, \boldsymbol{w}) &= \sum_{v} \sum_{\boldsymbol{w}:\max(\boldsymbol{w})=w} \frac{1}{N!} \sum_{\xi \in \Xi} \pi(v, \xi(\boldsymbol{w})) \\
&= \frac{1}{N!} \sum_{\xi \in \Xi} \sum_{v} \sum_{\xi(\boldsymbol{w}):\max(\xi(\boldsymbol{w}))=w} \pi(v, \xi(\boldsymbol{w})) \\
&= \frac{1}{N!} N! H(w) \\
&= H(w)
\end{aligned}$$

as needed. □

We can now show the proof for claim 8.

Proof. We start by showing that $Q^{BCEM}(H) \subseteq Q^{BCE}(H)$. let $\mu \in Q^{BCEM}$ and choose a $p(w, w^1, n^1, w^2, n^2, v) \in BCEM(H)$ that satisfies $\sum_{w,w^1,n^1,w^2,n^2} p(w, w^1, n^1, w^2, n^2, v) = \mu(v)$. We want to show that we can construct a symmetric BCE, π, which satisfies all $i \in \mathcal{N}$

$$\sum_{\substack{\pi: w_i = w, w^1 = \tilde{w}^1, w^2 = \tilde{w}^2, \\ n^1 = \tilde{n}^1, n^2 = \tilde{n}^2}} \pi(v, \boldsymbol{w}) = p(\tilde{w}, \tilde{w}^1, \tilde{n}^1, \tilde{w}^1, \tilde{n}^1, v) \tag{2.11}$$

Notice that such a BCE would clearly satisfy the obedience constraint and the data match constraint. Let

$$\Pi_i(\tilde{w}, \tilde{w}^1, \tilde{n}^1, \tilde{w}^2, \tilde{n}^2) = \Big\{ \boldsymbol{w} : \boldsymbol{w}_i = \tilde{w}, \boldsymbol{w}^1 = \tilde{w}^1, \boldsymbol{w}^2 = \tilde{w}^2,$$
$$|\{i : \boldsymbol{w}_i = \tilde{w}^1\} = \tilde{n}^1 |\{i : \boldsymbol{w}_i = \tilde{w}^2\} = \tilde{n}^2, \boldsymbol{w}_i \in \{w : p(w, \tilde{w}^1, \tilde{n}^1, \tilde{w}^2, \tilde{n}^2) > 0)\}, \Big\}$$

be the set of wage offers vectors in which firm i offers wage $\tilde{w}$, the other wage offers generate a distribution that satisfy the order statistics and includes only wage offers that are played with positive probability. Consider the case in which $w^1 > w^2$, $n^1 + n^2 = N$. Without loss of generality, we fix firm 1 and set for every joint probability of v, and $\boldsymbol{w} \in \Pi_1(w^1, w^1, n^1, w^2, n^2)$, the following

$$\pi(v, \boldsymbol{w}) = \frac{p(w^1, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^1-1}\binom{N-n^1}{n^2}}$$

and for every v, $\boldsymbol{w} \in \Pi_1(w^2, w^1, n^1, w^2, n^2)$ set

$$\pi(v, \boldsymbol{w}) = \frac{p(w^2, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^2-1}\binom{N-n^2}{n^1}}$$

The notice that that for each i and v and $w = w^1$ we have

$$\begin{aligned}
&\sum_{\boldsymbol{w}: w_i = w^1, w^1 = \tilde{w}^1 n^1 = \tilde{n}^1 w^2 = \tilde{w}^2 n^2 = \tilde{n}^2} \pi(v, \boldsymbol{w}) \\
&= \sum_{\boldsymbol{w}: w_1 = w^1, w_i = w^1, w^1 = \tilde{w}^1 n^1 = \tilde{n}^1 w^2 = \tilde{w}^2 n^2 = \tilde{n}^2} \pi(v, \boldsymbol{w}) + \sum_{\boldsymbol{w}: w_1 = w^2, w_i = w^1, w^1 = \tilde{w}^1 n^1 = \tilde{n}^1 w^2 = \tilde{w}^2 n^2 = \tilde{n}^2} \pi(v, \boldsymbol{w}) \\
&= \sum_{\boldsymbol{w}: w_1 = w^1, w_i = w^1, w^1 = \tilde{w}^1 n^1 = \tilde{n}^1 w^2 = \tilde{w}^2 n^2 = \tilde{n}^2} \frac{p(w^1, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^1-1}\binom{N-n^1}{n^2}} \\
&+ \sum_{\boldsymbol{w}: w_1 = w^2, w_i = w^1, w^1 = \tilde{w}^1 n^1 = \tilde{n}^1 w^2 = \tilde{w}^2 n^2 = \tilde{n}^2} \frac{p(w^2, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^2-1}\binom{N-n^2}{n^1}} \\
&= \sum_{\boldsymbol{w}: w_1 = w^1, w_i = w^1, w^1 = \tilde{w}^1 n^1 = \tilde{n}^1 w^2 = \tilde{w}^2 n^2 = \tilde{n}^2} \frac{p(w^1, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^1-1}\binom{N-n^1}{n^2}} \\
&+ \sum_{\boldsymbol{w}: w_1 = w^2, w_i = w^1, w^1 = \tilde{w}^1 n^1 = \tilde{n}^1 w^2 = \tilde{w}^2 n^2 = \tilde{n}^2} \frac{p(w^1, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^1-1}\binom{N-n^2}{n^1}} \\
&= \frac{p(w^1, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^1-1}\binom{N-n^1}{n^2}} \binom{N-1}{n^1-1} \binom{N-n^1}{n^2} = p(w^1, w^1, n^1, w^2, n^2, v)
\end{aligned}$$

where the third equality comes from constraint (2.8). An equivalent argument shows that this holds for every i and w^2. Next, consider the case in which $w^1 > w^1$ and $n^1 + n^2 < N$. Let $\mathcal{W} = \Big\{ w : w < w^2, p(w, w^1, n^1, w^2, n^2) > 0 \Big\}$ and define for each v and $\boldsymbol{w} \in \cup_{w \in \mathcal{W}} \Pi_1(w, w^1, n^1, w^2, n^2)$

$$\pi(v, \boldsymbol{w}) = \frac{p(w, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^1}\binom{N-1-n^1}{n^2}}$$

Set again, for every v, and $w \in \Pi_1(w^1, w^1, n^1, w^2, n^2)$, the following probability

$$\pi(v, \boldsymbol{w}) = \frac{p(w^1, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^1-1}\binom{N-n^1}{n^2}}$$

and for every v, $w \in \Pi_1(w^2, w^1, n^1, w^2, n^2)$ set

$$\pi(v, \boldsymbol{w}) = \frac{p(w^2, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^2-1}\binom{N-n^2}{n^1}}$$

Now, consider a firm i, making a wage offer $w \in \mathcal{W}$ and notice that

$$\begin{aligned}
&\sum_{\substack{\boldsymbol{w}: w_i = w, w^1 = \tilde{w}^1 n^1 = \tilde{n}^1 \\ w^2 = \tilde{w}^2 n^2 = \tilde{n}^2}} \pi(v, \boldsymbol{w}) = \\
&\sum_{\substack{\boldsymbol{w}: w_1 = w^1, w_i = w, w^1 = \tilde{w}^1, n^1 = \tilde{n}^1, \\ w^2 = \tilde{w}^2, n^2 = \tilde{n}^2}} \pi(v, \boldsymbol{w}) + \sum_{\substack{\boldsymbol{w}: w_1 = w^2, w_i = w, w^1 = \tilde{w}^1 n^1 = \tilde{n}^1 \\ w^2 = \tilde{w}^2 n^2 = \tilde{n}^2}} \pi(v, \boldsymbol{w}) \\
&+ \sum_{w \in \mathcal{W}} \sum_{\substack{\boldsymbol{w}: w_1 = w, w_i = w, w^1 = \tilde{w}^1 n^1 = \tilde{n}^1 \\ w^2 = \tilde{w}^2 n^2 = \tilde{n}^2}} \pi(v, \boldsymbol{w}) \\
&= \frac{p(w, w^1, n^1, w^2, n^2, v)}{\binom{N-1}{n^1}\binom{N-1-n^1}{n^2}} \binom{N-1}{n^1} \binom{N-1-n^1}{n^2} = p(w, w^1, n^1, w^2, n^2, v)
\end{aligned}$$

where again we used the (2.8) and (2.9). An analogous argument would show that for each i the marginalization of our constructed BCE satisfies 2.11 for each firm which plays w^1 and w^2.

Next, consider the case in which $w^1 = w^2$ and $n^1 = n^2 < N$. define again $\mathcal{W} = \Big\{ w : w < w^2, p(w, w^1, n^1, w^2, n^2) > 0 \Big\}$ and set for each v and each $w \in \Pi(w^1, w^1, n^1, w^1, n^1)$

$$\pi(v, \boldsymbol{w}) = \frac{p(w^1, w^1, n^1, w^1, n^1, v)}{\binom{N-1}{n^1-1}}$$

similarly for each $\boldsymbol{w} \in \cup_{w \in \mathcal{W}} \Pi_1(w, w^1, n^1, w^2, n^2)$, and for each v define

$$\pi(v, \boldsymbol{w}) = \frac{p(w, w^1, n^1, w^1, n^1, v)}{\binom{N-1}{n^1}}$$

A similar argument to the previous ones would show that marginalizing over these distribution satisfy 2.11. Finally, for the case in which $w^1 = w^2$ and $n^1 = n^2 = N$ define

$$\pi(v, \boldsymbol{w}) = p(w^1, w^1, N, w^1, N, v)$$

which clearly satisfy 2.11. Notice that by construction π satisfies data match and the obedience constraints and therefore $\pi \in BCE(H)$. Further notice that by construction $\sum_{\boldsymbol{w}} \pi(v, \boldsymbol{w}) = \mu(v)$ and therefore $\mu \in Q^{BCE}(H)$ which shows $Q^{BCEM}(H) \subseteq Q^{BCE}(H)$

Next, we turn to show that $Q^{BCE}(H) \subset Q^{BCEM}(H)$. The argument are similar to the argument made to show the reverse direction, but we keep the proof here for completion. Fix $\mu \in Q^{BCE}$. Let $Q^{BCESYM}(H) = \{\mu : \exists \pi \in BCE(H), \sum_w \pi(v, w) = \mu(v),$ and π symmetric $\}$. By claim 10 we know that $\mu \in Q^{BCESYM}(H)$ we can then show that $Q^{BCESYM}(H) \subset Q^{BCEM}(H)$. Fix a symmetric BCE $\pi \in BCE(H)$ such that $\sum_w \pi(v,w) = \mu(v) \forall v \in \mathcal{V}$. We can construct a $p(w, w^1, n^1, w^2, n^2)$ by marginalizing over π for a specific player. i.e. we define

$$p(\tilde{w}, \tilde{w}^1, \tilde{n}^1, \tilde{w}^2, \tilde{n}^2, v) = \sum_{\boldsymbol{w}: w_1 = w, w^1 = \tilde{w}^1, n^1 = \tilde{n}^1, w^2 = \tilde{w}^2, n^2 = \tilde{n}^2} \pi(v, \boldsymbol{w})$$

Notice that this construction immediately satisfies (2.6) and (2.7) and that the marginal of $\sum_{\tilde{w}, \tilde{w}^1, \tilde{n}^1, \tilde{w}^2, \tilde{n}^2} p(\tilde{w}, \tilde{w}^1, \tilde{n}^1, \tilde{w}^2, \tilde{n}^2, v) = \mu(v)$. To conclude the proof we need to show that $p(w, w^1, n^1, w^2, n^2)$ satisfies (2.8) - (2.10). To see that 2.8 is satisfied, let $X(w^1, n^1, x^2, n^2) \subset \{w^1, w^2, \underline{\mathrm{w}}\}^N$ be the set of vectors indicating which firm make a wage offer w^1, which make wage offer w^2 and who makes lower wage offer $\underline{\mathrm{w}} < w^2$, such that each vector satisfy $|\{i : x_i = w^1\}| = n^1$, $|\{i : x_i = w^2\}| = n^2$ and $|\{i : x_i = \underline{\mathrm{w}}\}| = N - n^1 - n^2$. Consider the case where $w^1 > w^2$ and $n^1 + n^2 = N$. Notice that due to symmetry we have that for each $x \in X(w^1, n^1, x^2, n^2)$ we have that $\sum_{\substack{\boldsymbol{w}: w_i = w^1 \forall i : x_i = w^1 \\ w_i = w^2 \forall i : x_i = w^2}} \pi(v, \boldsymbol{w}) = c$, where c is a constant. Then, notice that

$$p(w^1, w^1, n^1, w^2, n^2, v) = \sum_{\substack{\boldsymbol{w}: w_1 = w^1, \boldsymbol{w}^1 = w^1, \boldsymbol{n}^1 = n^1, \\ \boldsymbol{w}^2 = w^2, \boldsymbol{n}^2 = n^2}} = \binom{N-1}{n^1-1}\binom{N-n^1}{n^2} c$$

$$p(w^2, w^1, n^1, w^2, n^2, v) = \sum_{\substack{\boldsymbol{w}: w_1 = w^2, \boldsymbol{w}^1 = w^1, \boldsymbol{n}^1 = n^1, \\ \boldsymbol{w}^2 = w^2, \boldsymbol{n}^2 = n^2}} = \binom{N-1}{n^2-1}\binom{N-n^2}{n^1} c$$

which together implies (2.8). Similar line of arguments can show that this 2.9 and 2.10 are satisfied as well. Therefore, we show $p(w, w^1, n^1, w^2, n^2, v) \in BCEM(H)$ and therefore $\mu \in Q^{BCEM}(H)$ which concludes the proof. $\square$

2.8.2 *Illustrative Example*

To get a better intuition on the information contained in the observed wage distribution and the obedience constraints, we consider a simple illustrative example, with only a single firm making wage offers to workers. We assume that the worker productivity distribution lies on the finite support $\mathcal{V} = \{5, 10, 15\}$ and that the firm also offers wages from a finite set of wage offers $\mathcal{W} = \{5, 10, 15\}$. The marginal-profit for the firm from hiring a worker of type v, at wage w is $v - w$. Finally, we assume that workers accept the job offer, at wage w, only if the offered wage is $w >= v - 5$. Workers with $v = 5$ are willing to work for the firm at any wage $w \in \mathcal{W}$.[12] Let $p(10)$ and $p(5)$ be the share of workers who earn 10 and 5 in the data.[13] Before extending a wage offer, the firm observes certain signals on the worker productivity, $t \in \mathcal{T}$, which is unobserved by the analyst. Therefore the firm's interim-expected profit, by offering a wage W, is given by $\pi(w) = E[\mathbb{1}\{w > v - 5\}(v - w)|t]$. Let $F(w|t)$ be the wage setting rule for the firm, given the observed signal t, then a BNE satisfies that is F, such that for each w with $F(w|t) > 0$ we have $\pi(w) \geq \pi(w'), \forall w' \in \mathcal{W}$.

Using theorem 1, we can consider the set of possible distributions of v, by looking for a distribution of v, which satisfies the obedience and the data-match constraint. Specifically, let $\mathrm{P}(v, w)$ be the joint probability of observing a wage offer, w and a worker with productivity v, then the obedience

12. This reservation wage assumption assures us that the firm has an incentive to make wage offers higher than 5

13. Notice that offering a wage of 15 is a dominated strategy, and therefore we don't expect to see workers with wage 15

constraint gives us the following four inequalities

$$\begin{aligned} \mathrm{P}(15,10) &\geq \mathrm{P}(10,10) + \mathrm{P}(5,10) \quad (10 \to 5) \\ \mathrm{P}(10,5) + \mathrm{P}(5,5) &\geq \mathrm{P}(15,5) \quad (5 \to 10) \\ \mathrm{P}(15,10) + \mathrm{P}(10,10) + \mathrm{P}(5,10) &\geq 0 \quad (10 \to 15) \\ \mathrm{P}(10,5) + \mathrm{P}(5,5) &\geq 0 \quad (5 \to 15) \end{aligned} \tag{2.12}$$

where only the first two constraints bind. Now, consider that we want to derive bounds on the first moment of the workers productivity distribution. let $\mathrm{P}(v|w)$ be the probability of the state being v, given that the agent received a signal w. Then, using Bayes rule we can re-write these constraints as

$$\begin{aligned} \mathrm{P}(15|10) &\geq \mathrm{P}(10|10) + \mathrm{P}(5|10) \\ \mathrm{P}(10|5) + \mathrm{P}(5|5) &\geq \mathrm{P}(15|5) \end{aligned}$$

To derive the upper bound we can solve for

$$\begin{aligned} \max_{\mu(v)} \bar{E[v]} &= 15 \times \mathrm{P}(15) + 10 \times \mathrm{P}(10) + 5 \times \mathrm{P}(5) \\ &= 15 \times (\mathrm{P}(15|5)p(5) + \mathrm{P}(15|10)p(10)) \\ &+ 10 \times (\mathrm{P}(10|5)p(5) + \mathrm{P}(10|10)p(10)) \\ &+ 5 \times (\mathrm{P}(5|5)p(5) + \mathrm{P}(5|10)p(10)) \end{aligned}$$

Given the obedience constraint above and $\sum_v \mathrm{P}(v|w) = 1$ for each w. Notice that in order to maximize the above expression, we want to push as much weight onto $\mathrm{P}(15|w)$. However, The second obedience constraint constrains us from doing so, while still having the firm bid 10. For the firm to bid

10, the probability of gaining positive profit must be larger then the probability of losing. Therefore, to solve the maximization problem, we can set P(15|10) = 1, P(15|5) = 0.5 and P(10|5) = 0.5, and get the following upper bound

$$\overline{E[v]} = 12.5p(5) + 15p(10) = 15 - 2.5p(5)$$

Using a similar line of reasoning, and the first obedience constraint will give us the lower bound

$$\underline{\mathrm{E(v)}} = 5p(5) + 10p(10) = 10 - 5p(5)$$

From these bounds we can see that the data shows that only a small share of workers is earning high wages, then the distribution of workers cannot have too much weight on high values. And similarly, if the share of workers earning low wages is small, then it must be that there is a large share of workers with high productivity. Figure 2.3 below plots the upper and lower bound as a function of the $p(5)$.

Finally, notice that in this example we consider only one firm. In the general model introduced in section 2.2, the worker reservation wage was set by the other firms. This implies that actions on the part of one firm could not induce a profitable deviation in other firms. For example, consider an extreme case, in which we observe that the wage distribution is a degenerate distribution with point mass on 10. This can only be result of an equilibrium where both firms know that state is 10 with certainty, and therefore Bertrand competition pushes prices to 10. On the other hand, in the single firm example $E[v] \in [10, 15]$. The intuition for this is that in the reservation wage example, the reservation wage does not “react optimally" to the firms actions, and therefore, the set of possible outcomes is large. On the other hand, in the competitive environment, the firm can't only take into consideration the value of the worker but also needs to consider what the other firms will be willing

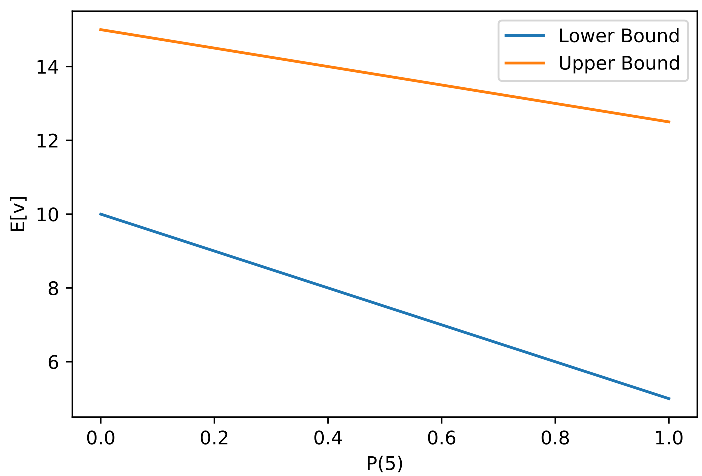

Figure 2.3: Upper and Lower bounds on the mean worker productivity in the single firm game

to offer to the worker.

References

Aigner, D. J. and Cain, G. G. (1977). Statistical theories of discrimination in labor markets. *ILR Rview*, 30(2):175–187.

Altonji, J. G. and Blank, R. M. (1999). Race and gender in the labor market. *Handbook of Labor Economics*, Chapter 48:3143–3259.

Altonji, J. G. and Pierret, C. R. (2001). Employer learning and statistical discrimination. *The Quarterly Journal of Economics*, 116(1):313–350.

Arrow, K. J. (1973). The theory of discrimination. *O. Ashenfelter and A. Rees (eds.), Discrimination in Labor Markets.*

Bergemann, D., Brooks, B., and Morris, S. (2017). First-price auctions with general information structures. *Econometrica*, 85(1):107–43.

Bergemann, D., Brooks, B., and Morris, S. (2021a). Counterfactuals with latent information. *Working Paper*.

Bergemann, D., Brooks, B., and Morris, S. (2021b). Search, information, and prices. *Journal of Political Economy*, 129(8):2275–2319.

Bergemann, D. and Morris, S. (2013). Robust predictions in games with incomplete information. *Econometrica*, 81(4):1251–1308.

Bergemann, D. and Morris, S. (2016). Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522.

Bergemann, D. and Morris, S. (2019a). Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95.

Bergemann, D. and Morris, S. (2019b). Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95.

Black, D. A. (1995). Discrimination in an equilibrium search model. *Journal of Labor Economics*, 12(2):309–24.

Blau, F. D. and Kahn., L. M. (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature*, 55(3):281–302.

Caldwell, S. and Danieli, O. (2021). Outside options in the labor market. *Working Paper*.

Chambers, C. P. and Echenique, F. (2021). A characterisation of 'phelpsian' statistical discrimination. *The Economic Journal*, 131(637):2018–2032.

Coate, S. and Loury, G. C. (1993). Will affirmative-action policies eliminate negative stereotypes? *The American Economic Review*, 83(5):1220–40.

Fang, Z., Santos, A., Shaikh, A., and Torgovitsky, A. (2020). Inference for large-scale linear systems with known coefficients. *SSRN Electronic Journal*.

Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review*, 104(4):1091–1119.

Gualdani, C. and Sinha, S. (2019). Identification and inference in discrete choice models with imperfect information. *arXiv preprint arXiv:1911.04529*.

Heckman, J. J., Lochner, L. J., and Todd, P. E. (2005). Earnings functions, rates of return and treatment effects: The mincer equation and beyond. *Working Paper Series. National Bureau of Economic Research*.

Lange, F. (2007). The speed of employer learning. *Journal of Labor Economics*, 25(1):1–35.

Lundberg, S. J. and Startz, R. (1983). Private discrimination and social intervention in competitive labor market. *The American Economic Review*, 73(3):340–47.

Magnolfi, L. and Roncoroni, C. (2017). Estimation of discrete games with weak assumptions on information. *Working Paper*.

Milgrom, P. and Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 1089–1122.

Mincer, J. (1958). Investment in human capital and personal income distribution. *Journal of Political Economy*, 66(4):281–302.

Phelps, E. S. (1972). The statistical theory of racism and sexism. *The american economic review*, 62(4):659–661.

Rubinstein, Y. and Weiss, Y. (2006). Chapter 1 post schooling wage growth: Investment, search and learning. *Handbook of the Economics of Education*, 11(2):1–67.

Shapley, L. S. and Shubik, M. (1971). The assignment game i: The core. *International Journal of Game Theory*, 1(1):30–111.

Spence, M. (1978). Job market signaling. *In Uncertainty in economics*, pages 281–306.

Syrgkanis, V., Tamer, E., and Ziani, J. (2021). Inference on auctions with weak assumptions on information. *arXiv preprint arXiv:1710.03830*.

CHAPTER 3

ON THE INTERPRETATION OF THE INTERGENERATIONAL ELASTICITY AND THE RANK-RANK COEFFICIENTS FOR CROSS COUNTRY COMPARISON

3.1 Introduction

Numerous studies have examined the relationship between parental income and child income. Two prominent methods for summarizing the joint distribution of these incomes are the Intergenerational Elasticity (IGE) coefficient and the Rank-Rank coefficient (Mogstad and Torsvik (2023)). This paper explores how these measures summarize the joint income distribution and their subsequent connections to the underlying mechanisms that link parental and child income.

Let I_c and I_p denote child and parent income, respectively. The IGE coefficient is the slope coefficient obtained by regressing the logarithm of child income on the logarithm of parent income as follows:

$$\log I_c = \alpha_{IGE} + \beta_{IGE} \log I_p + \epsilon. \tag{3.1}$$

This regression coefficient captures the persistence between child log income and the parent log income, with higher values indicating stronger persistence.[1] A popular alternative to this method is the Rank-Rank regression, which assesses the correlation between parent and child ranks within their respective income distributions. Assuming a continuous income distribution for both parents and children, let $R_c = F_c(I_c)$ and $R_p = F_p(I_p)$ represent the parent and child ranks in their respective income distributions, where $F_c(x)$ and $F_p(x)$ are the cumulative distribution functions of child and

1. In many cases, the level of intergenerational mobility is reported using (1-β_{IGE})

parental income, respectively. Researchers then measure the Rank-Rank relationship by estimating the following regression:

$$R_c = \alpha_r + \beta_r R_p + \varepsilon. \tag{3.2}$$

The regression slope coefficient quantifies how the child position in the income distribution relates to their parent position in the corresponding income distribution.

The IGE coefficient has been extensively employed in empirical studies to describe intergenerational persistence, dating back to the 1980s (Becker and Tomes (1986), Atkinson (1980)). However, the Rank-Rank coefficient has gained popularity more recently, after Chetty et al. (2014b) applied it to measure social mobility over time in the United States. While both coefficients are used to describe intergenerational mobility, each conveys distinct information about the joint distribution of parental and child income. As demonstrated below, the IGE provides a weighted average of the expected change in child logarithmic income in relation to a change in parent logarithmic income.[2] Consequently, the IGE coefficient is influenced by both the marginal distributions and the dependency structure between parental and child income. In contrast, the Rank-Rank coefficient measures positional mobility across generations, only summarizing the copula while isolating the dependency structure between the incomes and disregarding changes in marginal distributions (Deutscher and Mazumder (2023), Mogstad and Torsvik (2023), Aloni and Krill (2017)). From a practical perspective, the Rank-Rank coefficient has shown to more robust to sample restrictions (Chetty et al. (2014a), Chetty et al. (2014b), Dahl and DeLeire (2008)). In some countries (although not all; Bratberg et al. (2017), Acciari et al. (2022)), the Rank-Rank relation between parental and child income is almost

2. Mitnik and Grusky (2020) illustrates that the IGE can be considered as the elasticity of the conditional geometric mean, i.e., the expected percentage change in the geometric mean of the child's income with respect to the percentage change in the parental income.

perfectly linear. On the other hand, the conditional expectation function, $E[\log I_c | \log I_p]$, demonstrates significant nonlinearity (Chetty et al. (2014a), Deutscher and Mazumder (2023)). Moreover, the Rank-Rank coefficient allows researchers to include individuals with no income. This could be important since, as observed by Chetty et al. (2014a), the IGE demonstrates significant sensitivity to the substitution of zeros with ones or $1,000$s.

This paper examines the challenges that are inherent to using IGE and Rank-Rank coefficients for cross-country mobility comparisons. We express these coefficients as weighted averages of causal factors affecting intergenerational mobility using Yitzhaki's theorem (Yitzhaki (1996)), demonstrating that these coefficients assign varying weights across the parental income distribution. This helps to explain certain properties that were shown in the existing literature. We further explore how the parental income distribution influences the IGE and Rank-Rank coefficients, complicating cross-country comparisons, particularly when mobility occurs in different segments of the parental income distribution in each country. A related study (Maasoumi et al. (2022)) also employs Yitzhaki's theorem, framing the IGE coefficient weighting scheme as a special case within a broader class of intergenerational mobility measures that captures different preference relations over income distributions. The authors show that the IGE coefficient corresponds to a specific case of a preference relations that places higher weight on the mobility of wealthier households. In contrast, our study focuses on interpreting the coefficients as a weighted average of the underlying causal mechanisms and examines how these interpretations are important for cross-country comparison.

3.2 Decomposing the IGE coefficient

We begin by examining the β_{IGE} coefficient. Let us assume that (I_c, I_p) are i.i.d, $E[|\log I_c|], E[|\log I_p|] < \infty$, and $E[\log I_c | \log I_p = t]$ exists and is differentiable for all t. According to Yitzhaki's theorem, we can express β_{IGE} as a weighted average of the derivative of the conditional expectations:

$$\beta_{IGE} = \frac{\text{Cov}(\log I_c, \log I_p)}{\text{Var}(\log I_p)} = \int_{-\infty}^{\infty} \frac{\partial E[\log I_c | \log I_p = t]}{\partial t} w(t) dt,$$

where

$$w(t) = \frac{E\left[\log I_p - \mu_{I_p} | \log I_p > t\right] \text{P}(\log I_p > t)}{\text{Var}(\log I_p)}, \quad \int_{-\infty}^{\infty} w(t)dt = 1, \quad \mu_{I_p} = E[\log I_p].$$

We can then interpret the IGE coefficient as a summary statistic of the underlying function $E[\log I_c | \log I_p]$, where the weights depend on the distribution of parental income. Specifically, these weights are maximized at $E[\log I_p]$ and approach zero at the boundary of the support (Yitzhaki (1996), Heckman et al. (2006)). Thus, β_{IGE} assigns higher weight to households with the mean parental log income[3] and lower weights to households at the top and bottom of the parental log income distribution.

The fact that the IGE coefficient assigns lower weights to households at the extremes may be concerning in cases in which a significant portion of mobility occurs for children from very poor or very rich families. This can potentially occur as a result of policies aimed at reducing poverty or simply through regression to the mean. The fact that the weights depend on the underlying parental log income distribution implies that comparisons of the IGE coefficients that are cross-country or over time can be difficult to interpret. For instance, without knowing the exact weights, differences between two countries may simply arise from differences in the weighting schemes used by the IGE coefficient, even if the conditional expectation function $E[\log I_c | \log I_p]$ is the same across both countries.

Notably, the fact that the IGE coefficient assigns higher weights to mobility around the mean

3. Note that this is generally not the same as families with mean income.

may explain why the IGE is considered sensitive to sample definitions and restrictions. Some sample restrictions, such as excluding households with zero income or those with very high income, can significantly impact the mean of the distribution. As a result, households that receive higher weights change and the IGE coefficient also changes.

To better understand how the IGE coefficient relates to the Rank-Rank coefficient and underlying income elasticity,[4] we aim to decompose the integrand into the expected parent-child income elasticity and additional correlative mechanisms. Let the causal model governing child income be given as follows:

$$I_c = h(I_p, u), \tag{3.3}$$

where u represents other unobserved factors that affect child income. Let $\epsilon_{I_c,I_p}(u) = \frac{\partial \log I_c}{\partial \log I_p}$ be the elasticity of child income with respect to parent income for given unobserved factors, u, evaluated at I_p. Let $I_p(t) = exp(t)$ denote the inverse of $\log I_p$, where $\log I_p = t$. We can then rewrite the integrand as follows:

$$\begin{aligned}\frac{\partial E[\log I_c|\log I_p = t]}{\partial t} &= \int_{-\infty}^{\infty} \frac{\partial \log h\left(I_p\left(t\right),u\right) P(u|\log I_p = t)}{\partial t} du \\ &= \underbrace{E\left[\epsilon_{I_c,I_p(t)}(u)\middle|\log I_p = t\right]}_{\text{Causal IGE}} + \underbrace{\int_{-\infty}^{\infty} \log h\left(I_p\left(t\right),u\right) \frac{\partial P(u|\log I_p = t)}{\partial t} du}_{\text{Other Factors}},\end{aligned} \tag{3.4}$$

where the second equality follows from the product rule. The first component captures the conditional expected causal IGE, while the second component captures how changes in income are associated with

4. As noted by Mitnik and Grusky (2020), the IGE coefficient does not actually provide information about the parent-child income elasticity, which is evident in our setup, as $\frac{\partial E[\log I_c|\log I_p]}{\partial \log I_p} \neq E\left[\frac{\partial \log I_c}{\partial \log I_p}|\log I_p\right]$

changes in other factors that affect income.[5] Therefore, β_{IGE} can be expressed as a summation of the weighted causal intergenerational elasticities (causal IGE) and an additional term that captures how parental income is correlated with other factors that affect child income. In most studies of intergenerational mobility, both terms are crucial as researchers are interested in measuring how parent income is associated with child income, through either the causal effect of parental income or the association between parental income and other factors such as neighborhood quality, quality of schools, inherited human capital, and peer effects.

3.3 Decomposing the Rank-Rank coefficient

We now turn our attention to the Rank-Rank coefficient. Using Yitzhaki's theorem once more, we have the following:

$$\beta_r = \frac{\mathrm{Cov}(R_c, R_p)}{\mathrm{Var}(R_p)} = \int_{t=0}^{1} \frac{\partial E[R_c|R_p = t]}{\partial t} w(t)dt,$$

where, using the fact that the rank distribution is uniform, the exact weighting scheme is as follows:

$$w(t) = \frac{12(1-t)t}{2}, \quad \int_0^1 w(t)dt = 1.$$

Comparing the weights of the IGE coefficient to the Rank–Rank coefficient, the Rank–Rank weights place most of the weight on households at the median of the parental income distribution. In contrast, the IGE assigns most of the weight to households closer to the mean of the distribution. In addition, weights decline symmetrically as we move further away from the median and toward the extremes. Thus, similar to the β_{IGE} coefficient, the Rank–Rank coefficient assigns lower weights

5. This decomposition of the β_{IGE} can be thought of as an omitted variable bias. In this case, bias is taken with respect to the Ordinary Least Squares weighted causal effects of log parental income, as implied in Yitzhaki's theorem.

to households closer to the top and bottom of the parental income distribution. Notably, since the median is usually less sensitive to changes in sample restrictions at the top and bottom of the distributions, this weighting scheme might explain why the Rank–Rank coefficient has been documented to be more robust for different sample restrictions (Dahl and DeLeire (2008), Chetty et al. (2014b)). Finally, compared with the IGE coefficient, the weights for cross-country comparisons are more consistent, assigning similar weights to households at the same rank of the income distribution. Note that, if the marginal distributions differ across countries, this implies that the Rank–Rank weighting scheme assigns different weights to households with the same income levels. Whether this is desirable depends on the researcher's questions and objectives.

As we did for the IGE coefficient, we can express the Rank-Rank coefficients in terms of the underlying parent-child income elasticities. Let ϵ_c and ϵ_p be the elasticities of rank with respect to income for the child and parents, respectively. Let $R_c\epsilon_c = R_c\frac{\partial R_c}{\partial I_c}\frac{I_c}{R_c}$ and $R_p\epsilon_p = R_p\frac{\partial R_p}{\partial I_p}\frac{I_p}{R_p}$ represent the semi-elasticities of rank with respect to income. These quantities measure how the rankings of parents and child change in response to a percentage variation in their respective incomes. We can then rewrite, with a slight abuse of notation, the integrand as follows:[6]

$$
\begin{aligned}
\frac{\partial \mathrm{E}[R_c|R_p=t]}{\partial t} &= \frac{\partial \mathrm{E}[F_c\left(h\left(F_p^{-1}(t),u\right)\right)|R_p=t]}{\partial t} \\
&= \int_{-\infty}^{\infty} \frac{\partial F_c\left(h\left(F_p^{-1}\left(t\right),u\right)\right)P(u|R_p=t)}{\partial t} du \\
&= \mathrm{E}\left[\frac{\partial R_c}{\partial h}\frac{\partial h}{\partial I_p}\frac{1}{\frac{\partial R_p}{\partial I_p}}\middle| R_p=t\right] + \int_{\infty}^{\infty} F_c(h(F_p^{-1}(t),u))\frac{\partial P(u|R_p=t)}{\partial t}du \\
&= \mathrm{E}\left[\frac{\partial I_c}{\partial I_p}\frac{\frac{\partial R_c}{\partial I_c}}{\frac{\partial R_p}{\partial I_p}}\frac{I_c}{I_c}\frac{I_p}{I_p}\frac{R_c}{R_c}\frac{R_p}{R_p}\middle| R_p=t\right] + \int_{\infty}^{\infty} F_c(h(F_p^{-1}(t),u))\frac{\partial P(u|R_p=t)}{\partial t}du \\
&= \underbrace{E\left[\frac{R_c}{R_p}\frac{\epsilon_c}{\epsilon_p}\epsilon_{I_c,I_p}(u)\middle| R_p=t\right]}_{\text{Re-Scaled Causal IGE}} + \underbrace{\int_{\infty}^{\infty} F_c(h\left(F_p^{-1}(t),u\right)\frac{\partial P(u|R_p=t)}{\partial t}du}_{\text{Other factors}},
\end{aligned}
\tag{3.5}
$$

where the third equality is due to the product rule and the chain rule. The fourth equality results from dividing and multiplying by parents and child income and ranks and the definition of the parents and child ranks.[7] The final equality follows from the definition of semi-elasticities. Expressing the integrand in this way reveals the similarities and differences between the IGE coefficient and the Rank-Rank coefficient. First, as child income cumulative distribution function is monotonic, similar to the log function, the effects of other factors on income have remained the same, except that we use child marginal income distribution to transform the income instead of log. Likewise, the Rank-Rank coefficient is also affected by the causal effects of the IGE, but now the IGE is multiplied by a "translation" term that converts the income elasticities to rank elasticities.

6. Maasoumi et al. (2022) expresses the Rank-Rank coefficient as a weighted average of $\frac{\partial \mathrm{E}[\log I_c|\log I_p=t]}{\partial t}$, with weights that are generally positive but do not necessarily sum to 1. In contrast, we express the Rank-Rank coefficient as a weighted average of $\frac{\partial \mathrm{E}[R_c|R_p=t]}{\partial t}$ with weights that sum to 1.

7. For the sake of clarity, we slightly abuse notation and denote $I_c = h\left(F_p^{-1}(t),u\right)$, $R_c = F_c\left(h\left(F_p^{-1}(t),u\right)\right)$, $I_p = F_p^{-1}(t)$, and $R_p = t$.

If we are using the Rank–Rank coefficient for cross-country comparisons, the decomposition we derived above explicitly demonstrates that the Rank–Rank coefficient is only useful for comparisons of positional mobility. However, It cannot speak to how similar or different the mechanisms driving this mobility are across countries.[8] For example, consider two countries with the same underlying causal mechanisms $h(I_p, u)$ and assume that $I_p \perp\!\!\!\perp u$, which implies that the second term is zero. If the parental income distributions differ across the two countries, the Rank-Rank coefficient would still be different for two reasons. The first reason is that, although the weighting scheme is the same for households with the same income rank, the regression weighting scheme weights households with the same income level differently. The second and more substantial reason is that the way that the causal mechanisms affect rank would differ between the two countries as the semi-elasticities are different in the causal IGE term in equation 3.5. Therefore, although we might motivate the use of the Rank-Rank coefficient as a means to abstract away from the marginals, we cannot avoid considering the marginals if we want to use the Rank-Rank coefficient to think about differences in the driving mechanisms of mobility between two countries.

3.4 Discussion

This paper employs Yitzhaki's theorem to express IGE and Rank-Rank coefficients as weighted averages of the causal mechanisms driving income and positional mobility. We demonstrate that interpreting cross-country comparisons using the IGE coefficient can be challenging due to the regression weighting scheme. Additionally, we establish that the Rank-Rank coefficient is readily interpretable only when researchers focus on positional mobility, without providing insights into the similarities or

8. In theory, the Rank–Rank coefficient can be more informative on causal mechanisms that operate directly from parent income rank to child income rank, bypassing income levels. We leave this observation for future research.

differences in the underlying mechanisms driving mobility across countries.

We highlight the potential drawbacks of using linear regression coefficients as summary statistics. Linear regression may be preferred in certain cases for its efficiency and stability, even with a small number of observations. However, it seems that in the context of intergenerational mobility comparisons, this is not always warranted. Recent research has shifted to using large administrative datasets that can provide precise estimates of the relation between parent and child income. Consequently, the practice of reporting regression coefficients over estimates from more flexible and transparent methods may not always be well justified.

References

Acciari, Paolo, Alberto Polo, and Giovanni L. Violante. And Yet it Moves: Intergenerational Mobility in Italy. *American Economic Journal: Applied Economics*, 14(3):118–163, 2022.

Aloni, Tslil, and Zeev Krill. Intergenerational Income Mobility in Israel-International and Sectoral Comparisons. *Ministry of Economics, Office of Chief Economist*, 2017.

Atkinson, Anthony B. On Intergenerational Income Mobility in Britain. *Journal of Post Keynesian Economics*, 3(2):194–218, 1980.

Becker, Gary S., and Nigel Tomes. Human Capital and the Rise and Fall of Families. *Journal of Labor Economics*, 4(3, Part 2):S1–S39, 1986.

Bratberg, Espen, Jonathan Davis, Bhashkar Mazumder, Martin Nybom, Daniel D. Schnitzlein, and Kjell Vaage. A Comparison of Intergenerational Mobility Curves in Germany, Norway, Sweden, and the US. *The Scandinavian Journal of Economics*, 119(1):72–101, 2017.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States. *The Quarterly Journal of Economics*, 129(4):1553–1623, 2014.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, Emmanuel Saez, and Nicholas Turner. Is the United States Still a Land of Opportunity? Recent Trends in Intergenerational Mobility. *American Economic Review*, 104(5):141–147, 2014.

Dahl, Molly W., and Thomas DeLeire. The Association between Children's Earnings and Fathers' Lifetime Earnings: Estimates Using Administrative Data. *University of Wisconsin-Madison, Institute for Research on Poverty*, Madison, WI, USA, 2008.

Deutscher, Nathan, and Bhashkar Mazumder. Measuring intergenerational income mobility: A synthesis of approaches. *Journal of Economic Literature*, 61(3):988–1036, 2023.

Heckman, James J., Sergio Urzua, and Edward Vytlacil. Understanding Instrumental Variables in Models with Essential Heterogeneity. *The Review of Economics and Statistics*, 88(3):389–432, 2006.

Maasoumi, Esfandiar, Le Wang, and Daiqiang Zhang. Generalized Intergenerational Mobility Regressions. *Working paper, Emory University*, 2022.

Mitnik, Pablo A., and David B. Grusky. The Intergenerational Elasticity of What? The Case for Redefining the Workhorse Measure of Economic Mobility. *Sociological Methodology*, 50(1):47–95, 2020.

Mogstad, Magne, and Gaute Torsvik. Family Background, Neighborhoods, and Intergenerational Mobility. In Shelly Lundberg and Alessandra Voena (Eds.), *Handbook of the Economics of the Family*, Volume 1, Issue 1, pp. 327–387. Elsevier Science & Technology, 2023.

Yitzhaki, Shlomo. On Using Linear Regressions in Welfare Economics. *Journal of Business & Economic Statistics*, 14(4):478–486, 1996.